

Veritas™ Resiliency Platform 10.0 User Guide

Veritas™ Resiliency Platform User Guide

Last updated: 2022-04-22

Document version: Document version: 10.0 Rev 0

Legal Notice

Copyright © 2022 Veritas Technologies LLC. All rights reserved.

Veritas, the Veritas Logo, Veritas InfoScale, and NetBackup are trademarks or registered trademarks of Veritas Technologies LLC or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

This product may contain third-party software for which Veritas is required to provide attribution to the third party ("Third-Party Programs"). Some of the Third-Party Programs are available under open source or free software licenses. The License Agreement accompanying the Software does not alter any rights or obligations you may have under those open source or free software licenses. Refer to the third-party legal notices document accompanying this Veritas product or available at:

<https://www.veritas.com/about/legal/license-agreements>

The product described in this document is distributed under licenses restricting its use, copying, distribution, and decompilation/reverse engineering. No part of this document may be reproduced in any form by any means without prior written authorization of Veritas Technologies LLC and its licensors, if any.

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

The Licensed Software and Documentation are deemed to be commercial computer software as defined in FAR 12.212 and subject to restricted rights as defined in FAR Section 52.227-19 "Commercial Computer Software - Restricted Rights" and DFARS 227.7202, et seq. "Commercial Computer Software and Commercial Computer Software Documentation," as applicable, and any successor regulations, whether delivered by Veritas as on premises or hosted services. Any use, modification, reproduction release, performance, display or disclosure of the Licensed Software and Documentation by the U.S. Government shall be solely in accordance with the terms of this Agreement.

Veritas Technologies LLC
2625 Augustine Drive
Santa Clara, CA 95054

<http://www.veritas.com>

Technical Support

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

<https://www.veritas.com/support>

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

<https://sort.veritas.com/documents>

Documentation feedback

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

vrpdocs@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Section 1	Google Cloud Platform	10
Chapter 1	Recovering VMware virtual machines to Google Cloud Platform	11
	About recovering virtual machines to Google Cloud Platform	364
	Plan your environment	364
	Deploy and configure the virtual appliances	365
	Downloading the Veritas Resiliency Platform virtual appliances	366
	About deploying the Resiliency Platform virtual appliances	368
	Deploying the virtual appliances in AWS through AWS Marketplace	374
	Deploying the virtual appliances in AWS using OVA files	389
	Deploying the virtual appliances in Azure using PowerShell script	396
	Deploying the virtual appliances in Azure Stack using PowerShell script	401
	Deploying the virtual appliances in Azure through Azure Marketplace	405
	Deploy virtual appliances in Azure Stack using Azure Stack Marketplace	412
	Deploying the virtual appliances in vCloud	413
	Deploying the virtual appliances in Orange Recovery Engine	415
	Deploying the virtual appliance through VMware vSphere Client	419
	Deploying the virtual appliance through Hyper-V Manager	420
	Deploying the virtual appliances in Google Cloud Platform (GCP) through GCP Marketplace	421
	Deploying the virtual appliances in Google Cloud Platform using OVA files	429
	About configuring the Resiliency Platform components	437
	Set up the resiliency domain	453
	Getting started with a new configuration	453
	Adding an IMS	456
	Adding a Replication Gateway	111

Adding Google Cloud Platform data center	458
Managing user authentication and permissions	463
Managing settings for alerts and notifications and miscellaneous product settings	483
Add asset infrastructure	495
Adding VMware virtualization servers	167
Preparing host for replication	496
Infrastructure Pairing	508
About network objects	509
Network pairs for recovering virtual machines to Google Cloud Platform (GCP)	512
Creating network pairs between source and target data centers	514
Create resiliency groups	515
Configuring a resiliency group for basic monitoring	529
Managing virtual machines for remote recovery (DR) to Google Cloud Platform	517
Volume type selection options	519
Customize panel for Google Cloud Platform	520
Network customization options	523
About manual intervention	525
Advanced features	526
About virtual business services	526
About resiliency groups with assets	528
Configuring a resiliency group for basic monitoring	529
About evacuation plan	530
Perform remote recovery operations	532
Performing the rehearsal operation for virtual machines	532
Performing cleanup rehearsal for virtual machines	535
Migrating a resiliency group	536
Performing the resync operation for virtual machines	537
Recovering a resiliency group using replication-based recovery	539
Monitor assets	542
About risks	542
Predefined risks in Resiliency Platform	544
About reports	576
Scheduling a report	577
Running a report	579
Viewing reports	580
Managing a running activity	582
Miscellaneous references	583
About klish	584

Klish menu options for Resiliency Manager	586
Klish menu options for IMS	598
Klish menu options for Replication Gateway	612
About applying updates to Resiliency Platform	631
Virtual appliance security features	667

Chapter 2

Recovering Hyper-V virtual machines to Google Cloud Platform	363
About recovering virtual machines to Google Cloud Platform	364
Plan your environment	364
Deploy and configure the virtual appliances	365
Downloading the Veritas Resiliency Platform virtual appliances	366
About deploying the Resiliency Platform virtual appliances	368
Deploying the virtual appliances in AWS through AWS Marketplace	374
Deploying the virtual appliances in AWS using OVA files	389
Deploying the virtual appliances in Azure using PowerShell script	396
Deploying the virtual appliances in Azure Stack using PowerShell script	401
Deploying the virtual appliances in Azure through Azure Marketplace	405
Deploy virtual appliances in Azure Stack using Azure Stack Marketplace	412
Deploying the virtual appliances in vCloud	413
Deploying the virtual appliances in Orange Recovery Engine	415
Deploying the virtual appliance through VMware vSphere Client	419
Deploying the virtual appliance through Hyper-V Manager	420
Deploying the virtual appliances in Google Cloud Platform (GCP) through GCP Marketplace	421
Deploying the virtual appliances in Google Cloud Platform using OVA files	429
About configuring the Resiliency Platform components	437
Set up the resiliency domain	453
Getting started with a new configuration	453
Adding an IMS	456
Adding Google Cloud Platform data center	458
Managing user authentication and permissions	463
Managing settings for alerts and notifications and miscellaneous product settings	483

Add asset infrastructure	495
Adding Hyper-V virtualization servers	495
Preparing host for replication	496
Infrastructure Pairing	508
About network objects	509
Network pairs for recovering virtual machines to Google Cloud Platform (GCP)	512
Creating network pairs between source and target data centers	514
Create resiliency groups	515
Configuring a resiliency group for basic monitoring	529
Managing virtual machines for remote recovery (DR) to Google Cloud Platform	517
Volume type selection options	519
Customize panel for Google Cloud Platform	520
Network customization options	523
About manual intervention	525
Advanced features	526
About virtual business services	526
About resiliency groups with assets	528
Configuring a resiliency group for basic monitoring	529
About evacuation plan	530
Perform remote recovery operations	532
Performing the rehearsal operation for virtual machines	532
Performing cleanup rehearsal for virtual machines	535
Migrating a resiliency group	536
Performing the resync operation for virtual machines	537
Recovering a resiliency group using replication-based recovery	539
Monitor assets	542
About risks	542
Predefined risks in Resiliency Platform	544
About reports	576
Scheduling a report	577
Running a report	579
Viewing reports	580
Managing a running activity	582
Miscellaneous references	583
About klish	584
Klish menu options for Resiliency Manager	586
Klish menu options for IMS	598
Klish menu options for Replication Gateway	612
About applying updates to Resiliency Platform	631

	Virtual appliance security features	667
Chapter 3	Recovery to cloud data center	670
	Recovering VMware virtual machines to AWS	671
	Recovering Hyper-V virtual machines to AWS	675
	Recovering virtual machines from VMware to AWS using NetBackup Image Sharing	746
	Recovering VMware virtual machines to Azure	683
	Recovering Hyper-V virtual machines to Azure	687
	Recovering virtual machines from Azure / Azure Stack to Azure / Azure Stack	691
	Recovering VMware virtual machines to vCloud Director	696
	Recovering Hyper-V virtual machines to vCloud Director	700
	Recovering VMware virtual machines to vCloud Director without adding vCenter server	704
	Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server	708
	Recovering virtual machines from vCloud Director to vCloud Director	712
	Recovering VMware virtual machines to Orange Recovery Engine	716
	Recovering physical machines to AWS using Resiliency Platform Data Mover	719
	Recovering physical machines to vCloud Director using Resiliency Platform Data Mover	723
	Recovering physical machines to Orange Recovery Engine using Resiliency Platform Data Mover	727
	Recovering physical machines to Azure using Resiliency Platform Data Mover	730
Chapter 4	Recovery to on-premises data center	735
	Recovering physical machines to VMware virtual machines on an on-premises data center using Resiliency Platform Data Mover	735
	Recovering VMware virtual machines to on-premises data center using Resiliency Platform Data Mover	739
	Recovering VMware virtual machines from VMware to VMware using NetBackup	742
	Recovering virtual machines from VMware to AWS using NetBackup Image Sharing	746
	Recovering VMware virtual machines using third-party replication technology	750

Recovering Hyper-V virtual machines using third-party replication technology	753
Recovering Applications using third-party replication technology	756
Recovering InfoScale applications	759
Index	763
Glossary	767

Google Cloud Platform

- [Chapter 1. Recovering VMware virtual machines to Google Cloud Platform](#)
- [Chapter 2. Recovering Hyper-V virtual machines to Google Cloud Platform](#)

Recovering VMware virtual machines to Google Cloud Platform

This chapter includes the following topics:

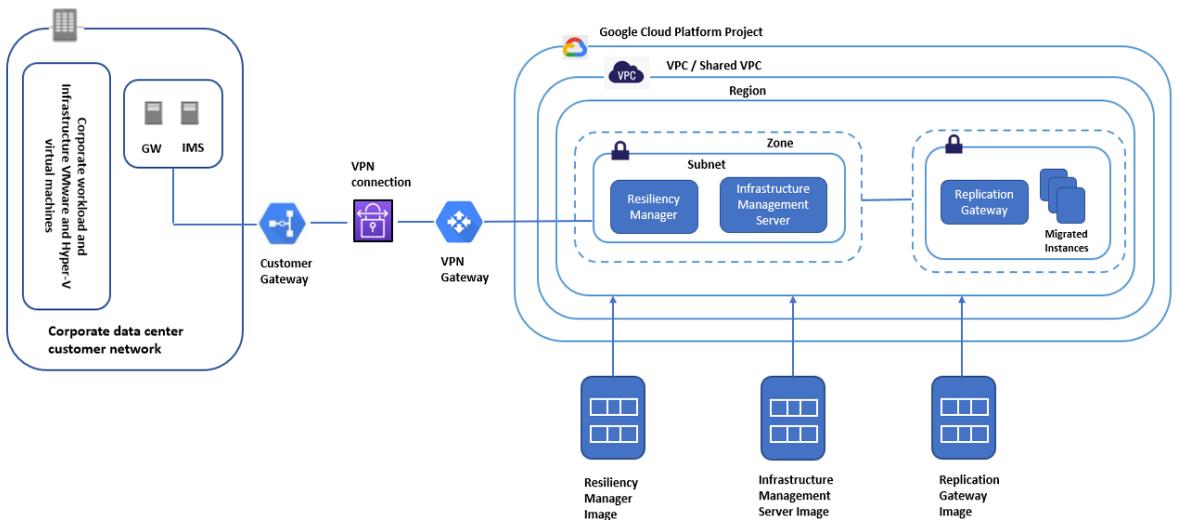
- [About recovering virtual machines to Google Cloud Platform](#)
- [Plan your environment](#)
- [Deploy and configure the virtual appliances](#)
- [Set up the resiliency domain](#)
- [Add asset infrastructure](#)
- [Infrastructure Pairing](#)
- [Create resiliency groups](#)
- [Advanced features](#)
- [Perform remote recovery operations](#)
- [Monitor assets](#)
- [Miscellaneous references](#)

About recovering virtual machines to Google Cloud Platform

Using Veritas Resiliency Platform, you can configure and protect your VMware virtual machines for recovery to Google Cloud Platform using the Resiliency Platform Data Mover. Following are the capabilities supported for this use case:

- Operations like Rehearsal, Recover, Migrate (and Migrate back), bulk-recovery are enabled when target data center is Google Cloud Platform.
- Network customization feature is added where you can enable or disable IP customization.
- Support for Shared VPC, which is a unique feature in Google Cloud Platform.
- Support for Regional disks, in addition to Zonal disks.
- Support for Customer Managed Encryption Keys.

Figure 1-1 Overview of deployment Infrastructure for recovery to Google Cloud Platform



Plan your environment

This topic explains about the way you are planning to configure the Resiliency Platform components in your environment. Refer to the **Overview and Planning**

Guide on SORT to know about the product, its components, features, and capabilities. Refer to the **Release Notes** on SORT for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.

How to find Resiliency Platform documents on SORT:

- 1 Navigate to [SORT home](#).
- 2 In the **Knowledge Base** tab, click on **Documents** link.
- 3 Select **Resiliency Platform** from the Product list. You can view all the links of the released versions.
- 4 Click on the **Unix or Windows** link to view all the guides.

Deploy and configure the virtual appliances

Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the Google Cloud Platform cloud data center as well as in the premises data center.

Refer to the following topics:

Deploy and configure virtual appliances in Google Cloud Platform :

To deploy the virtual appliances in the Google Cloud Platform region, refer the below sequence:

- See [“Downloading the Veritas Resiliency Platform virtual appliances”](#) on page 366.
- See [“About deploying the Resiliency Platform virtual appliances”](#) on page 368.
- See [“Deploying the virtual appliances in Google Cloud Platform \(GCP\) through GCP Marketplace”](#) on page 421.
- See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Deploy and configure virtual appliances in premise data center:

To deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center,

- See [“Deploying the virtual appliance through VMware vSphere Client”](#) on page 419. Deploying the virtual appliance through VMware vSphere Client

To configure the virtual appliances as Veritas Resiliency Platform components:

- See [“Prerequisites for configuring Resiliency Platform components”](#) on page 438.

- See “[About configuring the Resiliency Platform components](#)” on page 437.

Downloading the Veritas Resiliency Platform virtual appliances

You can download a licensed copy of the Veritas Resiliency Platform virtual appliances from [MyVeritas portal](#).

You can download the files for deploying virtual appliances. The virtual appliances are available in two formats: in the form of Open Virtualization Archive (OVA) files, or in the form of zip files. The .zip files contain the Virtual Hard disk (VHD) image file using which you can deploy the virtual appliances.

Table 1-1 Filenames for Veritas Resiliency Platform 10.0

Component	Filenames
Resiliency Manager	Veritas_Resiliency_Platform_RM_VMware_Virtual_Appliance_10.0.0.0_IE.ova
IMS	Veritas_Resiliency_Platform_IMS_VMware_Virtual_Appliance_10.0.0.0_IE.ova
Minimum hardware version for Resiliency Platform appliances	13 (ESXi 6.5 and above)
Resiliency Platform Data Mover (Up to hardware version 12)	Veritas_DataMover_VMware_Virtual_Appliance_10.0.0.0_IE.ova
Resiliency Platform Data Mover (From hardware version 14)	Veritas_DataMover_VMware_Virtual_Appliance_67_support_10.0.0.0_IE.ova

To download the Resiliency Platform virtual appliances:

- 1 Log in to MyVeritas portal:
<https://my.veritas.com>
- 2 Select **Licensing** tab, select the account and the entitlement ID that you want to use for downloading the Resiliency Platform virtual appliance.
- 3 In the list of products, click **Download** button next to Resiliency Platform.
- 4 Select the files that you want to download.

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

go.veritas.com/try-vrp

[Downloading the virtual appliances for Google Cloud Platform](#)

Downloading the virtual appliances for Google Cloud Platform

Below are the files which can be downloaded to be deployed on source and target data centers. The files are available in two formats: in the form of Open Virtualization Archive (OVA) files, or in the form of zip files. These .zip files contain the Virtual Hard disk (VHD) image file using which you can deploy the virtual appliances.

Table 1-2 Recovery of VMware virtual machines to Google Cloud Platform

Data Center	Component	Filenames
Source	IMS	Veritas_Resiliency_Platform_IMS_VMware_Virtual_Appliance_10.0.0.0_IE.ova
	Resiliency Platform Data Mover	Veritas_DataMover_VMware_Virtual_Appliance_10.0.0.0_IE.ova
Target	Resiliency Manager	Veritas_Resiliency_Platform_RM_VMware_Virtual_Appliance_10.0.0.0_IE.ova
	IMS	Veritas_Resiliency_Platform_IMS_VMware_Virtual_Appliance_10.0.0.0_IE.ova
	Resiliency Platform Data Mover	Veritas_DataMover_VMware_Virtual_Appliance_10.0.0.0_IE.ova

Table 1-3 Recovery of Hyper-V virtual machines to Google Cloud Platform

	Component	Filenames
Source	IMS	Veritas_Resiliency_Platform_IMS_Hyper-V_Virtual_Appliance_10.0.0.0_IE.zip
	Resiliency Platform Data Mover	Veritas_DataMover_Hyper-V_Virtual_Appliance_10.0.0.0_IE.zip

Table 1-3 Recovery of Hyper-V virtual machines to Google Cloud Platform
(continued)

	Component	Filenames
Target	Resiliency Manager	Veritas_Resiliency_Platform_RM_Hyper-V_Virtual_Appliance_10.0.0.0_IE.ova
	IMS	Veritas_Resiliency_Platform_IMS_Hyper-V_Virtual_Appliance_10.0.0.0_IE.ova
	Resiliency Platform Data Mover	Veritas_DataMover_Hyper-V_Virtual_Appliance_10.0.0.0_IE.ova

About deploying the Resiliency Platform virtual appliances

Veritas Resiliency Platform is deployed as a virtual appliance. A virtual appliance is a virtual machine image consisting of a pre-configured operating system environment with a software application installed on it. This virtual machine image can be deployed on a hypervisor. Once the Resiliency Platform virtual appliance gets deployed, you are required to configure the Resiliency Platform component through the product bootstrap.

Note: There is no sequence required for deploying and configuring the Resiliency Platform components. You can deploy and configure the components in any sequence on source as well as target data centers.

Following is the list of considerations for deploying the virtual appliances:

- For recovery to premises data center, you typically deploy and configure at least one Resiliency Manager and one Infrastructure Management Server (IMS) in the production data center and at least one Resiliency Manager and one Infrastructure Management Server (IMS) in the recovery data center.
- In case you plan to use Resiliency Platform Data Mover for recovery of your assets to premises data center, you need to deploy at least one Replication Gateway in the production data center and one Replication Gateway in the recovery data center.
- For recovery to cloud data center, you typically deploy and configure at least one Infrastructure Management Server (IMS) and one Replication Gateway in the production data center and one Resiliency Manager, one IMS, and one Replication Gateway in the recovery data center.

- The Replication Gateway on the production data center must have access to the ESX servers for the production virtual machines to be replicated. The Replication Gateway on the recovery data center must have access to the storage/ or compute or disk services of the target platform.

Resiliency Manager and IMS virtual appliance are shipped with single disk; similar to the Replication Gateway virtual appliance. While deploying the virtual appliances, it is prompted to attach a new empty disk of the required size. The RM disk size should be minimum 100 GB and IMS disk size should 40 GB.

While upgrading the virtual appliances, it will be prompted to attach the existing data disk of the previous version virtual appliances.

In the cloud environment (AWS, Azure and Google Cloud Platform), Marketplace offerings are available for upgrading the Veritas Resiliency Platform virtual appliances.

Refer to the topic **Deployment workflows** See [“Deployment workflows”](#) on page 370.

Based on the virtualization technology in your environment, choose any one of the following methods to deploy the virtual appliances in the on-premises data center:

Table 1-4 Deploying components in the on-premises data center

Virtualization technology	Steps to deploy the components
Hyper-V	See “Deploying the virtual appliance through Hyper-V Manager” on page 420.
VMware	See “Deploying the virtual appliance through VMware vSphere Client” on page 419.

Based on your cloud data center, choose any one of the following methods to deploy the virtual appliances in the cloud data center:

Table 1-5 Deploying components in the cloud data center

Cloud data center	Steps to deploy the components
Google Cloud Platform	See “Deploying the virtual appliances in Google Cloud Platform (GCP) through GCP Marketplace” on page 421. See “Deploying the virtual appliances in Google Cloud Platform using OVA files” on page 429.

Table 1-5 Deploying components in the cloud data center (*continued*)

Cloud data center	Steps to deploy the components
AWS	See “Deploying the virtual appliances in AWS through AWS Marketplace” on page 374. See “Deploying the virtual appliances in AWS using OVA files” on page 389.
vCloud Director	See “Deploying the virtual appliances in vCloud” on page 413.
Azure	See “Deploying the virtual appliances in Azure through Azure Marketplace” on page 405. See “Deploying the virtual appliances in Azure using PowerShell script” on page 396.
Azure Stack	See “Deploying the virtual appliances in Azure Stack using PowerShell script” on page 401. See “Deploy virtual appliances in Azure Stack using Azure Stack Marketplace” on page 412.
Orange Recovery Engine	See “Deploying the virtual appliances in Orange Recovery Engine” on page 415.

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

Deployment workflows

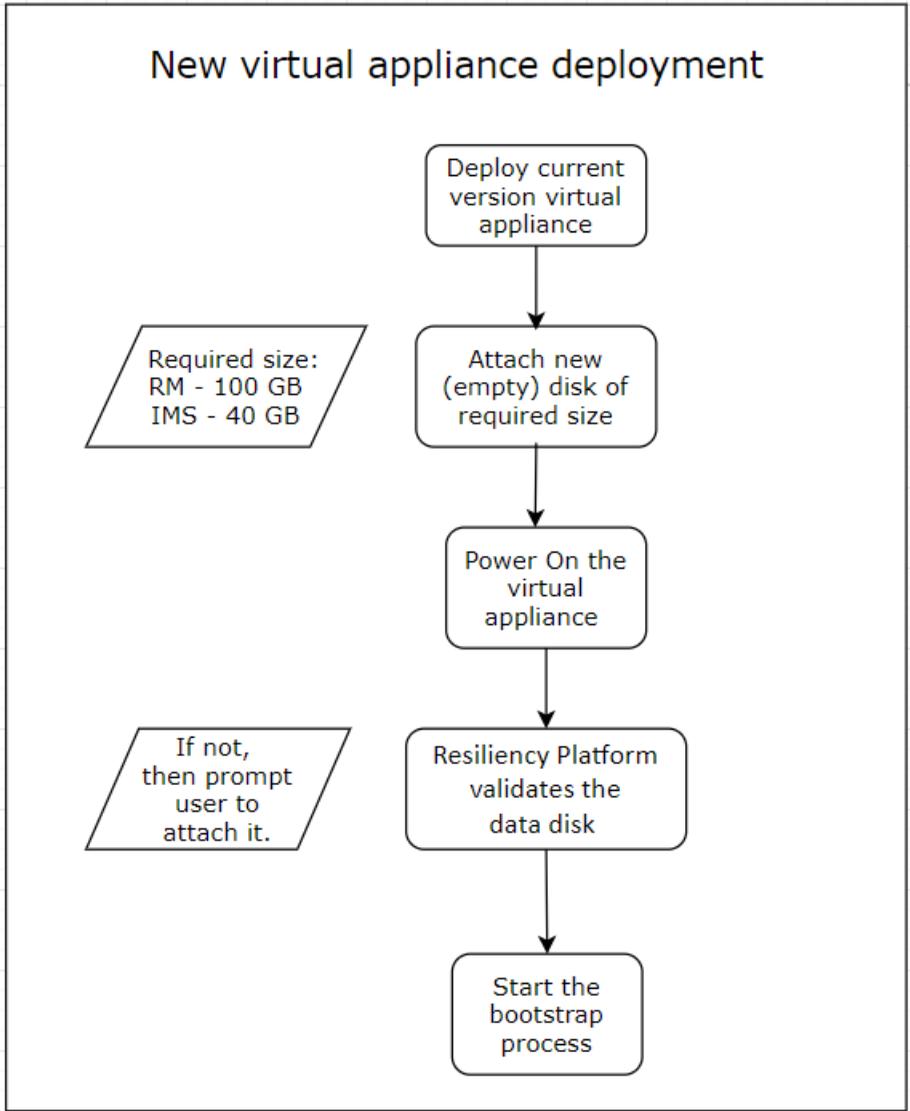
The new approach for upgrading the Resiliency Platform (Resiliency Manager and IMS) involves saving configuration to the data disk of the previous version virtual appliance and then attaching the data disk to a new, freshly deployed virtual appliance.

The Resiliency Platform virtual appliances would be created along with the OS disk. It is required to attach an existing data disk while upgrading the Resiliency Platform virtual appliances. While deploying the new virtual appliances, attach a new empty disk.

This approach is similar to that of deploying the Replication Gateway virtual appliance. It is now applicable to Resiliency Manager and IMS too. Once the disk is attached, the workflow for both the Resiliency Platform upgrade and bootstrap process is same.

Below is the deployment workflow introduced for Resiliency Manager and IMS virtual appliance:

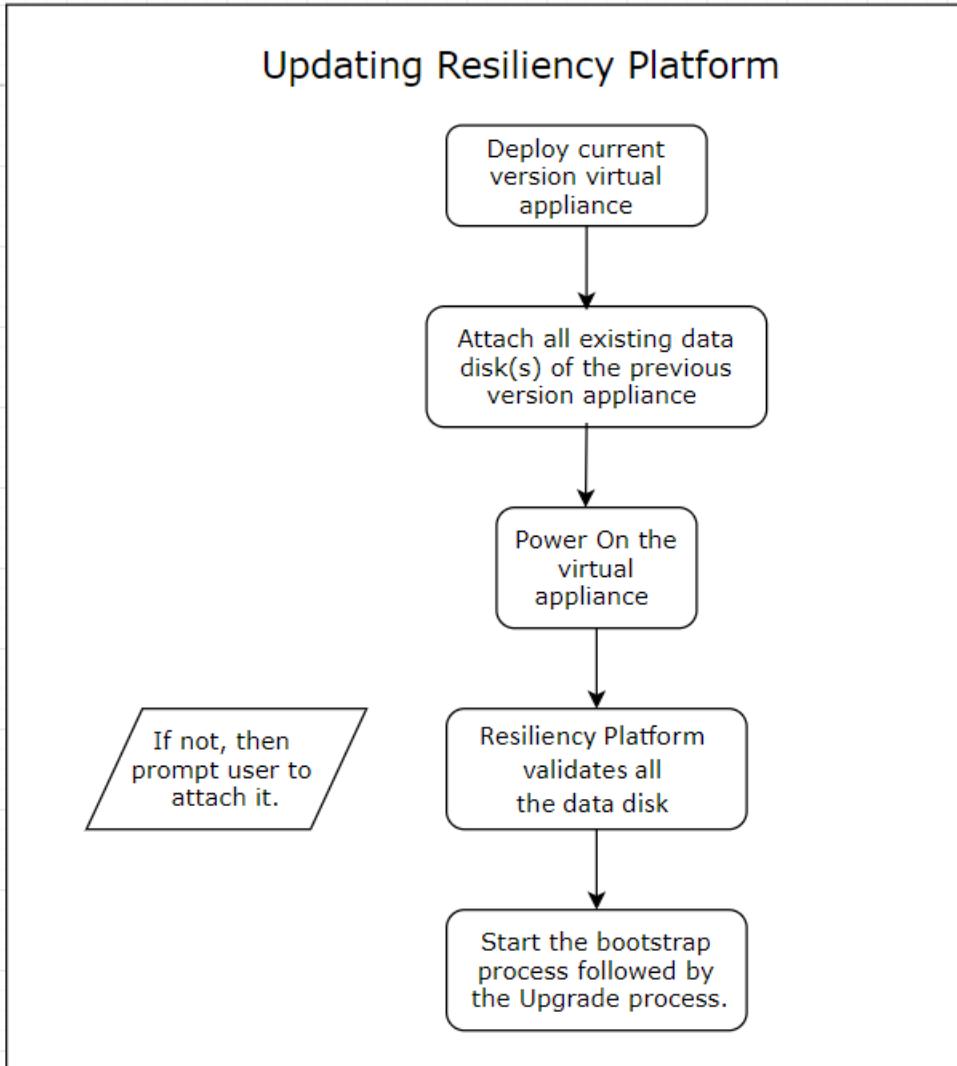
Figure 1-2 Deploying the new Resiliency Platform virtual appliances



Steps for deploying the new Resiliency Platform virtual appliances

- 1** Download the virtual appliances and deploy the current version virtual appliances.
- 2** After successfully deploying the appliances, attach an empty disk for the appliances. The required disk size for RM is minimum 100 GB and IMS is 40 GB.
- 3** Power on the virtual appliances.
- 4** Resiliency Platform validates the attached disk. Even after power on, the disk can be attached. You will be prompted to attach the disk if not attached.
- 5** Start the bootstrap process to further continue with the configuration.

Figure 1-3 Upgrading (update) the Resiliency Platform virtual appliances



Steps for upgrading Resiliency Platform virtual appliances

- 1 Download the virtual appliances and deploy the current version virtual appliances.
- 2 After successfully deploying the appliances, attach all the existing data disks of the previous version appliances.

- 3 Power on the virtual appliances.
- 4 Resiliency Platform validates all the attached data disk. Even after power on, the disk can be attached. You will be prompted to attach the data disk if not attached.
- 5 Start the bootstrap process followed by the upgrade process to further continue with the configuration.

Note: Ignore any disk or filesystem related messages given by the operating system while executing the Resiliency Platform virtual appliance upgrade or deployment bootstrap process.

Deploying the virtual appliances in AWS through AWS Marketplace

Veritas Resiliency Platform enables you to deploy the virtual appliances in AWS through AWS Marketplace using CloudFormation templates. There are six offerings available for deploying the virtual appliances using CloudFormation templates:

- **Veritas Resiliency Platform Express Install:** Installs Resiliency Manager, IMS, and Replication Gateway appliances in AWS. This template also provides options to install Veritas Data Gateway appliance in AWS.
- **Veritas Resiliency Platform Gateway Install:** Installs an additional Replication Gateway appliance in AWS.
- **Veritas Resiliency Platform Resiliency Manager Install:** Installs an additional Resiliency Manager appliance in AWS.
- **Veritas Resiliency Platform Infrastructure Management Install :** Installs the Infrastructure Management server in AWS.
- **Veritas Resiliency Platform Resiliency Manager Upgrade :** Upgrades the Resiliency Manager appliance in AWS.
- **Veritas Resiliency Platform Infrastructure Management Upgrade :** Upgrades the Infrastructure Management server in AWS.

To deploy the virtual appliances in AWS using CloudFormation templates

- 1 Prerequisites
Create Amazon Virtual Private Cloud (VPC):
- 2 Go to the **AWS Marketplace** and locate **Veritas Resiliency Platform** product.
- 3 Select the fulfillment option from the options available: **Express Install**, **Gateway Install**, **Resiliency Manager**, or **Infrastructure Management Server**.

- 4 AWS Marketplace lets you launch the selected option through CloudFormation Templates interface. You are redirected to the **Create Stack** page of AWS CloudFormation Template. The template URL is pre-populated for you. Click **Next**.
- 5 On the next page, provide the values for the input fields:
See [“Providing inputs for Resiliency Platform CloudFormation Templates”](#) on page 381.
- 6 Click **Next** and review the additional options for your stack. Select the check box in the **Capabilities** section.
- 7 Click **Next** and review the summary displayed on the **Review** page.
- 8 Click **Create** to create the instances. This step creates the EC2 instances, staging volume for Replication Gateway appliance, and required Security Groups. This step also completes the bootstrap for all the appliances.
- 9 For multiple Network Interface card (NIC) instances, every NIC created is associated with a different security group.

- 10 For security reasons, the CloudFormation template disables the SSH communication to the instances. Once the stack deployment completes, you need to enable the SSH communication to the instances for setting the admin password. This password is required for logging in to the Resiliency Manager console. A password is also required for adding the IMS, Replication Gateway appliances to Resiliency Manager.

Modify the security groups of each instance to include the inbound SSH port (TCP port 22) for the required source IP or security group.

Note: The instances go through some automatic configuration steps immediately after deployment. If you connect to the instance via SSH and if the configuration is still in-progress, wait until all steps are completed and reconnect to the instance.

- 11 Log in to the Resiliency Manager console and setup the initial infrastructure through Getting Started wizard.

Note: For **Express install**, if creation of stack fails and rollback is enabled on failure, AWS takes care of all the clean-up required in this situation. However, to enable this clean-up, you need to manually delete the Elastic Network Interface (ENI) associated with the security group that is attached with **ZipFileUploaderLambda** function.

Note: Instance Metadata Service v2 (IMDSv2) is introduced by AWS as security enhancement over IMDSv1. While using AWS CloudFormation template, there is no way to disable the IMDSv1 option. So this limits the AWS Marketplace deployment to have IMDSv2 enabled. It can be changed later using AWS CLI.

See [“Recovering VMware virtual machines to AWS”](#) on page 671.

See [“Recovering Hyper-V virtual machines to AWS”](#) on page 675.

Prerequisites for deploying the virtual appliances in AWS

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

- Follow the documentation of AWS to create the required security groups. make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open for communication.

If you deploy Resiliency Platform components through AWS Marketplace, the required security groups are automatically created.

- Create a role named *vmimport* and grant the permissions required to import an image to the role. Follow the documentation of AWS to know about the permissions required to import an image.

The *vmimport* role should have the following KMS service permissions (along with the permissions mentioned in the AWS documentation):

- "kms:ReEncrypt"
 - "kms:GenerateDataKeyWithoutPlaintext"
 - "kms:DescribeKey"
 - "kms:CreateGrant"
 - "kms:Decrypt"
- Create Individual roles for Resiliency Manager, IMS, and Replication Gateway with certain permissions. These roles are used for authenticating the operations performed by the Resiliency Platform components in AWS.

See [“Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway”](#) on page 389.

If you deploy Resiliency Platform through AWS Marketplace, then the required role are automatically created through AWS CloudFormation template that the marketplace deployment uses.

- Ensure that there is direct communication between the premise network and the AWS network. It is recommended to use VPN for AWS environment.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. You may choose to restrict incoming internet access on these virtual appliances.
- Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.
- For significant cost benefits, it is recommended to buy Amazon EC2 reserved instances for the Resiliency Platform virtual appliances as the virtual appliances will be running continuously. While buying the reserved instances, it is also recommended to select the availability zones where the virtual appliances are to be deployed; this will ensure reserved capacity. It is important to choose the reserved instance types that matches the virtual appliances' instance types.

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

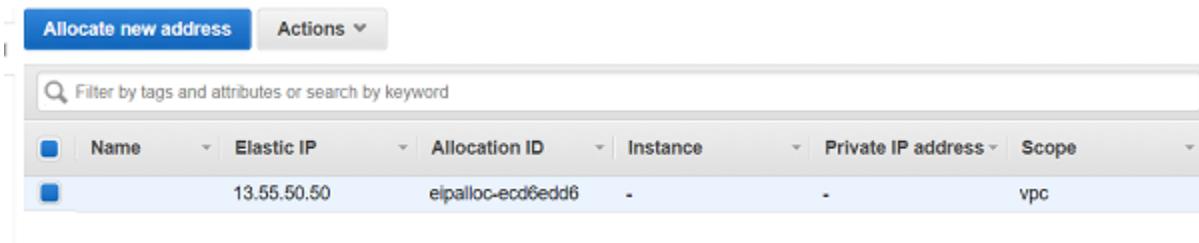
Configuring Amazon VPC for deployment using CloudFormation Templates

You need to configure Amazon Virtual Private Cloud (VPC) for deploying the Veritas Resiliency Platform components in AWS using CloudFormation Templates (CFT). For information on deploying the Veritas Resiliency Platform components in AWS using CFT:

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

To configure Amazon VPC

- 1 Access the VPC service in the AWS region where you want to deploy Resiliency Platform components, and go to the **Elastic IPs** tab.
- 2 Click on **Allocate new address** to allocate an Elastic IP. This Elastic IP will be required for the NAT gateway while creating the VPC.



- 3 In the **VPC Dashboard** tab, click on **Start VPC Wizard**.

- 4 In the **Select a VPC Configuration** step, select the **VPC with Public and Private Subnets** option.

Step 1: Select a VPC Configuration

VPC with a Single Public Subnet

VPC with Public and Private Subnets

VPC with Public and Private Subnets and Hardware VPN Access

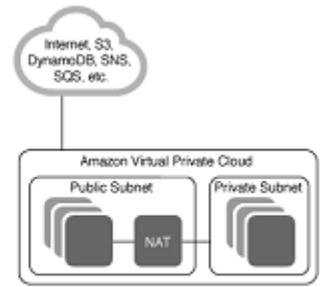
VPC with a Private Subnet Only and Hardware VPN Access

In addition to containing a public subnet, this configuration adds a private subnet whose instances are not addressable from the Internet. Instances in the private subnet can establish outbound connections to the Internet via the public subnet using Network Address Translation (NAT).

Creates:

A /16 network with two /24 subnets. Public subnet instances use Elastic IPs to access the Internet. Private subnet instances access the Internet via Network Address Translation (NAT). (Hourly charges for NAT devices apply.)

[Select](#)



The diagram illustrates an Amazon Virtual Private Cloud (VPC) setup. At the top, a cloud icon represents the Internet, containing services like S3, DynamoDB, SNS, and SQS. Below this, a box labeled 'Amazon Virtual Private Cloud' contains two subnets: 'Public Subnet' and 'Private Subnet'. A 'NAT' gateway is positioned between the two subnets, with lines indicating traffic flow from the private subnet through the NAT gateway to the public subnet, and then to the Internet.

- 5 In the **VPC with Public and Private Subnets** page, do the following:
 - Choose an appropriate CIDR block for the VPC and give it the desired name. Repeat the same steps for the public and private subnets.
 - For the NAT Gateway, select the Elastic IP that was allocated in step 2.
 - Ensure that the **Enable DNS hostnames** option is set to **Yes**.

Step 2: VPC with Public and Private Subnets

IPv4 CIDR block:* (65531 IP addresses available)

IPv6 CIDR block: No IPv6 CIDR Block
 Amazon provided IPv6 CIDR block

VPC name:

Public subnet's IPv4 CIDR:* (251 IP addresses available)

Availability Zone:*

Public subnet name:

Private subnet's IPv4 CIDR:* (251 IP addresses available)

Availability Zone:*

Private subnet name:

You can add more subnets after AWS creates the VPC.

Specify the details of your NAT gateway ([NAT gateway rates apply](#)).

Elastic IP Allocation ID:*

Service endpoints

Add Endpoint

Enable DNS hostnames:* Yes No

Hardware tenancy:*

6 Finish the wizard to create the VPC.

Providing inputs for Resiliency Platform CloudFormation Templates

You need to provide inputs for creating instances using CloudFormation templates (CFT). Some of the fields get auto populated with the default value, you can change the values if required. For rest of the parameters, you need to enter a valid value.

Table 1-6 EC2 Instance Configuration

Field	Description
Resiliency Manager Instance Name	Enter the name of the Resiliency Manager instance.
IMS Instance Name	Enter the name of the IMS instance.
Replication Gateway Instance Name	Enter the name of the Replication Gateway instance.
EC2 Instance Type	This field is auto-populated with the best fit instance type. You can change the default value, if required.
Key Pair for SSH access	Select the appropriate SSH key pair. While selecting the SSH key pair, ensure that you already have access to the pair. This is critical because SSH key pair is the only way to log in to the system and you cannot change the SSH key pair after the EC2 instance gets created.
Create 'ImportSnapshotRole' AWS IAM Role for VRP?	Resiliency Platform requires ImportSnapshotRole for recovery of assets to AWS. Select Yes to create the role. Make sure that you have iam: CreateRole permission. Select No if you already have ImportSnapshotRole role.
Staging Disk size for the Replication Gateway	Enter the size of the staging disk for Replication Gateway in the unit of GB. The minimum supported disk size of 50GB lets you protect up to 8 virtual machines and to protect more virtual machines an additional staging storage of 6 GB per virtual machine is required.

Table 1-6 EC2 Instance Configuration (*continued*)

Field	Description
Data Volume Type	Select the appropriate Volume Type for the data volume. This field has default value selected as 'gp2' volume type. You can change and select another as per your requirement.
Resiliency Manager Data volume IOPS (Required only when selected Volume Type is IO1 from IO1 and IO2)	This value will be used when selected Volume Type is IO1 only .IOPS value must be between 100 to 64000.
Infrastructure Management Server Data volume IOPS (Required only when selected Volume Type is IO1 from IO1 and IO2)	This value will be used when selected Volume Type is IO1 only. IOPS value must be between 100 to 64000.
Replication Gateway Staging Disk IOPS (Required only when selected Volume Type is IO1 from IO1 and IO2)	<p>This value will be used when selected DataVolumeType is IO1. The maximum ratio of provisioned IOPS to requested volume size (in GiB) is 50:1. For example, a 100 GiB volume can be provisioned with up to 5,000 IOPS. If appropriate IOPS value is not selected , the deployment fails.</p> <p>This value will be used when selected DataVolumeType is IO2. The maximum ratio of provisioned IOPS to requested volume size (in GiB) is 500:1. For example, a 100 GiB volume can be provisioned with up to 50,000 IOPS. If appropriate IOPS value is not selected , the deployment fails.</p>
KMS Key ARN for Staging Disk encryption	If you want to encrypt staging disks attached to the gateways using the KMS key, then select ARN of the respective KMS from the KMS Key drop-down. Else keep the default value as 'Not Encrypted'.

Table 1-7 Resiliency Manager Network Configuration

Field	Description
Network Interface to be used for communication with other Resiliency Managers	<p>Select the network interface that can be used to communicate with other Resiliency Managers in the resiliency domain.</p> <p>NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.</p>
Network Interface to be used for communication with Infrastructure Management Servers	<p>Select the network interface that can be used to communicate with Infrastructure Management Servers.</p> <p>NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.</p>
Network Interfaces to be used for accessing the User Interface	<p>Select the network interfaces that can be used to access the Resiliency Manager web user interface.</p> <p>NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.</p>
Network Interface to be used as the default gateway	<p>Select the network interface that can be used by default for outgoing communication.</p>
Resiliency Manager eth0 Subnet	<p>Select the subnet which can be connected to the eth0 network interface.</p>
Resiliency Manager eth1 Subnet	<p>Select the subnet which can be connected to the eth1 network interface.</p> <p>This input is only relevant if you have selected eth1 for one of the communication roles. You should provide a valid subnet always.</p>
Is the eth0 Network Interface behind NAT?	<p>If NAT settings are applicable for the communications you have selected for eth0 and if you wish to use NAT with eth0, select Yes.</p>
Resiliency Manager eth0 NAT Hostname (Optional)	<p>Provide the NAT hostname for eth0 if applicable</p>

Table 1-7 Resiliency Manager Network Configuration (*continued*)

Field	Description
Resiliency Manager eth0 NAT IP (Optional)	Provide the NAT IP for eth0 if applicable
Is the eth1 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth1 and if you wish to use NAT with eth1, select Yes.
Resiliency Manager eth1 NAT Hostname (Optional)	Provide the NAT hostname for eth1 if applicable
Resiliency Manager eth1 NAT IP (Optional)	Provide the NAT IP for eth1 if applicable

Table 1-8 Infrastructure Management Server Network Configuration

Field	Description
Network Interface to be used for communication with Resiliency Managers	Select the network interface that can be used to communicate with Resiliency Managers. NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.
Network Interface to be used for communication with Replication Gateways	Select the network interface that can be used to communicate with Replication Gateways. NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.
Network Interface to be used as the default gateway	Select the network interface that can be used by default for outgoing communication.
IMS eth0 Subnet	Select the subnet which can be connected to the eth0 network interface.
IMS eth1 Subnet	Select the subnet which can be connected to the eth1 network interface. This input is only relevant if you have selected eth1 for one of the communication roles. You should provide a valid subnet always.

Table 1-8 Infrastructure Management Server Network Configuration
(continued)

Field	Description
Is the eth0 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth0 and if you wish to use NAT with eth0, select Yes.
IMS eth0 NAT Hostname (Optional)	Provide the NAT hostname for eth0 if applicable
IMS eth0 NAT IP (Optional)	Provide the NAT IP for eth0 if applicable
Is the eth1 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth1 and if you wish to use NAT with eth1, select Yes.
IMS eth1 NAT Hostname (Optional)	Provide the NAT hostname for eth1 if applicable
IMS eth1 NAT IP (Optional)	Provide the NAT IP for eth1 if applicable

Table 1-9 Replication Gateway Network Configuration

Field	Description
Network Interface to be used for communication with the Infrastructure Management Server	Select the network interface that can be used to communicate with the Infrastructure Management Server. NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.
Network Interface to be used for communication with peer gateways	Select the network interface that can be used to communicate with peer gateways. NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.

Table 1-9 Replication Gateway Network Configuration (*continued*)

Field	Description
Network Interface to be used for communication with workload Virtual Machines	<p>Select the network interfaces that can be used for communication with the workload Virtual Machines.</p> <p>NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.</p>
Network Interface to be used as the default gateway	Select the network interface that can be used by default for outgoing communication.
Replication Gateway eth0 Subnet	Select the subnet which can be connected to the eth0 network interface.
Replication Gateway eth1 Subnet	<p>Select the subnet which can be connected to the eth1 network interface.</p> <p>This input is only relevant if you have selected eth1 for one of the communication roles. You should provide a valid subnet always.</p>
Replication Gateway eth2 Subnet	<p>Select the subnet which can be connected to the eth2 network interface.</p> <p>This input is only relevant if you have selected eth2 for one of the communication roles. You should provide a valid subnet always.</p>
Is the eth0 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth0 and if you wish to use NAT with eth0, select Yes.
Replication Gateway eth0 NAT Hostname (Optional)	Provide the NAT hostname for eth0 if applicable
Replication Gateway eth0 NAT IP (Optional)	Provide the NAT IP for eth0 if applicable
Is the eth1 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth1 and if you wish to use NAT with eth1, select Yes.
Replication Gateway eth1 NAT Hostname (Optional)	Provide the NAT hostname for eth1 if applicable

Table 1-9 Replication Gateway Network Configuration (*continued*)

Field	Description
Replication Gateway eth1 NAT IP (Optional)	Provide the NAT hostname for eth1 if applicable
Is the eth2 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth2 and if you wish to use NAT with eth2, select Yes.
Replication Gateway eth2 NAT Hostname (Optional)	Provide the NAT hostname for eth2 if applicable
Replication Gateway eth2 NAT IP (Optional)	Provide the NAT IP for eth2 if applicable

Table 1-10 Common Network Configuration

Field	Description
VPC ID	Ensure the following while selecting the VPC: <ul style="list-style-type: none"> Outgoing internet access is enabled from at least one private subnet from the VPC. VPC has a VPN configured with the network at the on- premises data center.

Table 1-11 Resiliency Platform Installation Parameters

Field	Description
NTP Server	FQDN or IP address of the NTP server to be used for the instances. In case of multiple values, enter the space-separated values.
TimeZone	Select the timezone for the instances.

Table 1-12 Data Gateway Deployment Information

Field	Description
Data Gateway deployment bucket name	Make sure that the deployment bucket pre-exists and is created in the local region where the CFT is being deployed.
Deploy Veritas Data Gateway?	Optional parameter You need to deploy Data Gateway only if you want to use Object Storage for replication.

Table 1-12 Data Gateway Deployment Information (*continued*)

Field	Description
Does the Data Gateway bucket already exist?	Select Yes if the Data Gateway already exists. If you want to re-use an existing Data Gateway Bucket, that bucket should also be local to the region in which the CFT is being deployed.
SNS Topic Protocol	Select a protocol from the list.
SNS Topic Endpoint	Provide an endpoint that is appropriate for the selected SNS topic protocol.

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Uninstalling Resiliency Platform components when deployed through AWS Marketplace

When you deploy Resiliency Platform components through AWS Marketplace, a stack gets created in AWS environment for each offering that you use. For example, if you use Express install, one stack gets created and all the Resiliency Platform components are deployed in that stack. If you use Resiliency Manager offering to install additional Resiliency Manager, another stack gets created. If you use Replication Gateway offering to install additional Replication Gateway, one more stack gets created.

If at a later point of time, you want to uninstall all the Resiliency Platform components that were deployed through AWS Marketplace, you need to delete all these stacks individually. Once you delete the stacks, all the resources that were created during deployment automatically get deleted.

To delete a stack in AWS

- 1 From the list of stacks in the AWS CloudFormation console, select the stack that you want to delete.
- 2 Click **Actions > Delete stack**.
- 3 When prompted, confirm that you want to delete the stack.

If you want to delete a stack created for **Express install** within 6 hours of creating the stack, you need to manually delete the Elastic Network Interface (ENI) associated with the security group that is attached with **ZipFileUploaderLambda** function.

Rest of the stack gets deleted automatically when you perform **Delete stack** operation.

Note: The NICs attached to Resiliency Platform virtual appliances that are deployed through the Marketplace will not be deleted automatically when the instances are deleted. They would need to be cleaned up manually after deletion.

Deploying the virtual appliances in AWS using OVA files

To know about virtual appliance deployment in Veritas Resiliency Platform:

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

Table 1-13 Overview of deployment process in AWS

Step	Action	Description
1	Ensure that the prerequisites for deploying virtual appliances in AWS are met.	See “Prerequisites for deploying the virtual appliances in AWS” on page 376.
2	Upload the OVA files to Amazon S3	See “Uploading the OVA file using web-based method” on page 393. See “Uploading the OVA file using command-line method” on page 394.
3	Create AMI using EC2	See “Creating Amazon Machine Image” on page 394.
4	Launch the instances of virtual appliances to deploy Resiliency Manager, Infrastructure Manager (IMS), and Replication Gateway	See “Launching the instances of virtual appliances” on page 395.

Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway

Following are the permissions required for the roles that you need to create for Resiliency Manager, IMS, and Replication Gateway for recovery to AWS data center.

Table 1-14 Permissions required for role for Resiliency Manager

Service name	Permission
ec2	ec2:DescribeVpcs ec2:DescribeAvailabilityZones
S3	s3:ListBucket

Table 1-15 Permissions required for role for IMS

Service name	Permission
ec2	

Table 1-15 Permissions required for role for IMS (*continued*)

Service name	Permission
	ec2:RunInstances
	ec2:CreateTags
	ec2:StartInstances
	ec2:DescribeInstanceStatus
	ec2:StopInstances
	ec2:TerminateInstances
	ec2:RegisterImage
	ec2:DescribeImageAttribute
	ec2:DescribeInstanceAttribute
	ec2:DeregisterImage
	ec2:DeleteVolume
	ec2:CreateVolume
	ec2:AttachVolume
	ec2:DetachVolume
	ec2:CreateSnapshot
	ec2>DeleteSnapshot
	ec2:ImportSnapshot
	ec2:DescribeImportSnapshotTasks
	ec2:CreateImage
	ec2:CreateNetworkInterface
	ec2>DeleteNetworkInterface
	ec2:AttachNetworkInterface
	ec2:ModifyNetworkInterfaceAttribute
	ec2:DescribeVpcs
	ec2:DescribeSubnets
	ec2:DescribeImages
	ec2:DescribeSecurityGroups
	ec2:DescribeNetworkInterfaces
	ec2:DescribeVolumes
	ec2:DescribeAvailabilityZones

Table 1-15 Permissions required for role for IMS (*continued*)

Service name	Permission
	ec2:DescribeSnapshots ec2:DescribeInstances ec2:GetEbsDefaultKmsKeyId ec2:GetEbsEncryptionByDefault ec2:DescribeInstanceTypes
execute-api	execute-api:Invoke
S3	s3:ListBucket s3:PutObject s3>DeleteObject s3:GetObject s3:GetBucketLocation
KMS	kms:ListAliases kms:ListKeys kms:DescribeKey

Table 1-16 Permissions required for role for Replication Gateway

Service name	Permission
execute-api	execute-api:Invoke
S3	s3:PutObject s3:GetObject s3:ListBucket

Uploading the OVA file using web-based method

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

To upload the OVA file using web-based method

- 1 Log in to the AWS console and go to **Services**.
- 2 Go to **S3**, and click **Create a bucket**.
- 3 Enter a name for the bucket and select the appropriate region. Click **Create**.

- 4 Once the bucket gets created, open the bucket and click **Upload**. Click **Add files** and then select the OVA file from your local disk.
- 5 Click **Start Upload**.

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Uploading the OVA file using command-line method

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

To upload the OVA file using command-line method

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

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

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

If you do not use the region option of the command, the bucket is created in the region specified in your configuration file.

- 3 Upload the OVA file by running the following command:

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

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

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Creating Amazon Machine Image

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

To create Amazon Machine Image

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

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

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

- 3 Run the following command to create an AMI:

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

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

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

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

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

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Launching the instances of virtual appliances

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

Instance MetaData Service v2 (IMDSv2) is introduced by AWS as security enhancement over IMDSv1. Instance MetaData Service v2 (IMDSv2) is introduced by AWS as security enhancement over IMDSv1. From version 4.0 of Resiliency Platform, while configuring resiliency group for disaster recovery, the wizard has an option to specify metadata access using **Enforce IMDSv2** option per virtual machine. If this option is true, the virtual machine when migrated to AWS, should use only IMDSv2 mechanism.

To launch the instances of virtual appliances

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

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

In the **Select an existing key pair or create a new key pair** wizard, you can choose an existing key pair, or create a new one. If you create a new key pair, ensure to click the **Download key pair** button and download. You will need the private key from this key pair to login as admin user for completing the bootstrap.

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Deploying the virtual appliances in Azure using PowerShell script

To know about virtual appliance deployment in Veritas Resiliency Platform:

You can deploy the Resiliency Platform virtual appliances using the following two files provided by Veritas along with the product files:

- A PowerShell script that handles the entire deployment of virtual appliances in Azure environment.
- A text file that has all the parameters required for deployment in Azure environment. You need to update the values assigned to the parameters in the text file. These values are used by the PowerShell script while deploying the virtual appliances.

To deploy the virtual appliances in Azure

- 1** Ensure that the prerequisites for deploying virtual appliances in Azure are met. See [“Prerequisites for deploying the virtual appliances in Azure and Azure Stack”](#) on page 403.
- 2** Log in to the Azure portal and create a general purpose storage account (and not account type as blob storage). Create a container with private access type under this storage account. Follow the documentation of Azure to create the storage account and container.
- 3** Create a static network interface for the virtual appliance. A static network interface in Azure ensures that the IP of the appliance does not change after reboot.
- 4** Download or copy the Azure deployment files on your local Windows system.

- 5 Update the values of all the parameters in the text file `VRPVSADeployInputs.txt` according to your environment.

Virtual Machine Appliance	Compatible Sizes	Recommended Sizes
Resiliency Manager	Standard_DS5_v2	Standard_D8s_v3
	Standard_DS13_v2	Standard_DS5_v2
	Standard_E8s_v3	
	Standard_D8s_v3	
	Standard_E8_v3	
	Standard_F16s_v2	
	Standard_F16s	
	Standard_A8m_v2	
Infrastructure Management Server	Standard_D13	
	Standard_F8s	Standard_F8s
	Standard_F8s_v2	Standard_F8s_v2
	Standard_D8s_v3	Standard_A8_v2
	Standard_DS4_v2	
	Standard_D4_v2	
	Standard_DS4	
	Standard_D4	
Replication Gateway	Standard_A8_v2	
	Standard_A4	
	Standard_F8s	Standard_F8s
	Standard_F8s_v2	Standard_F8s_v2
	Standard_D8s_v3	Standard_A8_v2
	Standard_DS4_v2	
	Standard_D4_v2	
	Standard_DS4	
Standard_D4		
Standard_A8_v2		
Standard_A4		

It is recommended that the Replication Gateway should be of virtual machine size which supports ultra disk and premium storage disk type.

- 6 Open the PowerShell command prompt and run the following command to log in to Azure:

```
Add-AzureRmAccount
```

- 7 Enter the absolute path and filename and then press the **Enter** key to run the Powershell script.

Prerequisites for deploying the virtual appliances in Azure and Azure Stack

Following are the prerequisites for deploying the Resiliency Platform virtual appliances in Azure. These prerequisites are applicable to deploy virtual appliances in Azure Stack too:

- Follow the documentation of Azure to create the required security groups. make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open.
- Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.
- Ensure that there is direct communication between the premise network and the Azure network. It is recommended to use VPN for Azure environment.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. You may choose to restrict incoming internet access on these virtual appliances.
- Ensure that PowerShell is installed on the Windows system from which you plan to run the Azure deployment script.
- Ensure that Azure PowerShell module is installed on the system.
See [“Installing Azure PowerShell module”](#) on page 399.

See [“Deploying the virtual appliances in Azure using PowerShell script”](#) on page 396.

Installing Azure PowerShell module

Installing Azure PowerShell from the PowerShell Gallery is the preferred method of installation.

To install Azure PowerShell module

- 1 You should have PowerShellGet module installed on your system.

<https://www.microsoft.com/en-us/download/details.aspx?id=51451>

- 2 You can verify if PowerShellGet module is properly installed on your system by executing the following command:

```
Get-Module PowerShellGet -list | Select-Object Name,Version,Path
```

If PowerShellGet module is properly installed on your system, you get the output of the above command similar to the following:

```
PowerShellGet 1.0.0.1 C:\Program  
Files\WindowsPowerShell\Modules\PowerShellGet\1.0.0.1\PowerShellGet.psd1
```

- 3 Once PowerShellGet module gets installed on your system, install the **Azure Resource Manager** modules by running the following command as an administrator:

- `Install-Module -Name Az`

- 4 You can check if Azure Resource Manager is properly installed by running any one of the following commands:

- `Add-AzAccount`

- `Get-AzResourceGroup`

See “[Prerequisites for deploying the virtual appliances in Azure and Azure Stack](#)” on page 403.

See “[Deploying the virtual appliances in AWS through AWS Marketplace](#)” on page 374.

Constraints when you deploy the virtual appliances in Azure

Following are a few limitations that apply when you deploy the Resiliency Platform virtual appliances in Azure:

- To prevent any security vulnerabilities, `Waagent` service is stopped 30 minutes after the bootstrapping completes on all the Resiliency Platform appliances. If required, you can start the service for next 30 minutes by using the following Klish command:

```
azure-waagent-service start
```

- You must not use `run-command` option from Azure portal on any of the Resiliency Platform appliances. Running `run-command` on the appliances may cause some functionality related issues for the product.

- Extensions are disabled on all the Resiliency Platform appliances.

Deploying the virtual appliances in Azure Stack using PowerShell script

To know about virtual appliances in Veritas Resiliency Platform:

You can deploy Resiliency Platform virtual appliances using the following two files provided by Veritas along with the product files:

- A PowerShell script that handles the entire deployment of virtual appliances in Azure Stack environment.
- A text file that has all the parameters required for deployment in Azure Stack environment. You need to update the values assigned to the parameters in the text file. These values are used by the PowerShell script while deploying the virtual appliances

To deploy the virtual appliances in Azure

- 1** Ensure that the prerequisites for deploying virtual appliances in Azure Stack are met.

See [“Prerequisites for deploying the virtual appliances in Azure and Azure Stack”](#) on page 403.
- 2** Log in to the Azure Stack portal and create a general purpose storage account (and not account type as blob storage). Create a container with private access type under this storage account. Follow the documentation of Azure to create the storage account and container.
- 3** Create a static network interface for the virtual appliance. A static network interface in Azure ensures that the IP of the appliance does not change after reboot.
- 4** Download or copy the Azure deployment files on your local Windows system.

- 5 Update the values of all the parameters in the text file `VRPVSADeployInputs.txt` according to your environment.

Virtual Machine Appliance	Compatible Sizes	Recommended Sizes
Resiliency Manager	Standard_DS5_v2	Standard_D8s_v3
	Standard_DS13_v2	Standard_DS5_v2
	Standard_E8s_v3	
	Standard_D8s_v3	
	Standard_E8_v3	
	Standard_F16s_v2	
	Standard_F16s	
	Standard_A8m_v2	
Infrastructure Management Server	Standard_D13	
	Standard_F8s	Standard_F8s
	Standard_F8s_v2	Standard_F8s_v2
	Standard_D8s_v3	Standard_A8_v2
	Standard_DS4_v2	
	Standard_D4_v2	
	Standard_DS4	
	Standard_D4	
Replication Gateway	Standard_AB_v2	
	Standard_A4	
	Standard_F8s	Standard_F8s
	Standard_F8s_v2	Standard_F8s_v2
	Standard_D8s_v3	Standard_A8_v2
	Standard_DS4_v2	
	Standard_D4_v2	
	Standard_DS4	
Standard_D4		
Standard_A8_v2		
Standard_A4		

It is recommended that the Replication Gateway should be of virtual machine size which support premium storage type.

- 6 Open the PowerShell command prompt and run the following command to log in to Azure:

To add a new Azure environment:

```
Add-AzEnvironment -Name <env_name> -ArmEndpoint <endpoint_url>
```

To connect to Azure environment:

```
Connect-AzAccount -Environment <env_name> -SubscriptionId  
<subscription_id>
```

- 7 Enter the absolute path and filename and then press the **Enter** key to run the PowerShell script.

Prerequisites for deploying the virtual appliances in Azure and Azure Stack

Following are the prerequisites for deploying the Resiliency Platform virtual appliances in Azure. These prerequisites are applicable to deploy virtual appliances in Azure Stack too:

- Follow the documentation of Azure to create the required security groups. make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open.
- Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.
- Ensure that there is direct communication between the premise network and the Azure network. It is recommended to use VPN for Azure environment.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. You may choose to restrict incoming internet access on these virtual appliances.
- Ensure that PowerShell is installed on the Windows system from which you plan to run the Azure deployment script.
- Ensure that Azure PowerShell module is installed on the system. See [“Installing Azure PowerShell module”](#) on page 399.

See [“Deploying the virtual appliances in Azure using PowerShell script”](#) on page 396.

Installing Azure Stack PowerShell module

Installation of PowerShell module depends on Azure stack version. You can refer Microsoft documentation. The below mentioned steps are for Azure Stack version 1908:

- Install Windows PowerShell 5.1.
- Install `PowerShellGet`: You need access to the PowerShell Gallery. The gallery is the central repository for PowerShell content. The `PowerShellGet` module contains cmdlets for discovering, installing, updating, and publishing PowerShell artifacts.
 - To install a package from the Gallery either execute the `Install-Module` or `Install-Script` cmdlet, depending on the package type.
Run: `Install-Module -Name PowershellGet -RequiredVersion 2.2.1`
 - Validate the PowerShell Gallery accessibility using the below commands:
`Import-Module -Name PowerShellGet -ErrorAction Stop`
`Import-Module -Name PackageManagement -ErrorAction Stop`
`Get-PSRepository -Name "PSGallery"`
 - If the repository isn't registered, open an elevated PowerShell session and run the following command:
`Register-PSRepository -Default Set-PSRepository -Name "PSGallery" -InstallationPolicy Trusted`

Install Azure Stack PowerShell

To install the Azure Stack PowerShell, perform the below steps:

- Install the `AzureRM.BootStrapper` module. Select **Yes** when prompted to install NuGet.
`Install-Module -Name AzureRM.BootStrapper`
- Install and import the API Version Profile required by Azure Stack into the current PowerShell session. Required version will vary with respect to Azure Stack Version
`Use-AzureRmProfile -Profile 2019-03-01-hybrid -Force Install-Module -Name AzureStack -RequiredVersion 1.7.2`

Confirm the installation

You can confirm the Azure Stack PowerShell installation by executing the below commands:

```
Get-Module -Name "Azure*" -ListAvailable  
Get-Module -Name "Azs*" -ListAvailable
```

See [“Deploying the virtual appliances in Azure Stack using PowerShell script”](#) on page 401.

Constraints when you deploy the virtual appliances in Azure Stack

Following are a few limitations that apply when you deploy the Resiliency Platform virtual appliances in Azure Stack:

- To prevent any security vulnerabilities, `Waagent` service is stopped 30 minutes after the bootstrapping completes on all the Resiliency Platform appliances. If required, you can start the service for next 30 minutes by using the following Klish command:

```
azure-waagent-service start
```
- You must not use `run-command` option from Azure portal on any of the Resiliency Platform appliances. Running `run-command` on the appliances may cause some functionality related issues for the product
- Extensions are disabled on all the Resiliency Platform appliances.

See [“Deploying the virtual appliances in Azure Stack using PowerShell script”](#) on page 401.

Deploying the virtual appliances in Azure through Azure Marketplace

To know about virtual appliance deployment in Veritas Resiliency Platform:

A few constraints are applied when you deploy Resiliency Platform in Azure:

See [“Constraints when you deploy the virtual appliances in Azure”](#) on page 400.

Veritas Resiliency Platform enables you to deploy the virtual appliances in Azure through Azure Marketplace using Azure Resource Manager (ARM) templates. To deploy the appliances search for the offering **Veritas Resiliency Platform on the Azure marketplace..**

The offering provide choice to create a set of Resiliency Manager, Infrastructure Management server (IMS) and Replication Gateway or to deploy these appliances individually.

To deploy the virtual appliances in Azure through Azure Marketplace

1 Prerequisites

Ensure that the prerequisites are met:

See [“Prerequisites for deploying the virtual appliances in Azure and Azure Stack”](#) on page 403.

2 Go to the Azure Marketplace and locate **Veritas Resiliency Platform** offering.

3 Select the offering. Azure Marketplace lets you launch the selected offering.

4 You are redirected to the **Deployment** page of the offering. Click **Create** to initiate the process.

5 On the next page, provide the values for the input fields:

See [“Providing inputs for deploying virtual appliances through Azure Marketplace”](#) on page 406.

6 Click **OK** and review the summary displayed in the **Summary** section.

7 Click **OK** to view and accept the terms and conditions.

8 Click **Create** to create the instances:

- This step creates selected virtual appliance(s) and completes the bootstrap of appliances with the provided inputs. It also creates the required network security groups and network interfaces. An additional staging disk is also created during Replication Gateway deployment.

9 For security reasons, the template disables the SSH communication to the instances. Once the deployment completes, you need to enable the SSH communication to the instances. The password is required for logging in to the Resiliency Manager console. The password is also required for adding the IMS, Replication Gateway to Resiliency Manager.

Modify the security group of each instance to include the inbound SSH port (TCP port 22) for the required source IP or security group.

10 Log in to the Resiliency Manager console and setup the initial infrastructure through Getting Started wizard.

Providing inputs for deploying virtual appliances through Azure Marketplace

You need to provide inputs for creating instances using Azure Resource Manager (ARM) templates. Some of the fields get auto populated with the default value, you can change the values if required. For rest of the parameters, you need to enter a valid value.

To create the appliance images, you need to provide inputs in following three sections:

- Basics [Table 1-17](#)
- Advanced [Table 1-18](#)
- Network [Table 1-19](#)

Figure 1-4 Basic settings for Azure deployment

The screenshot shows the 'Basics' step of an Azure deployment configuration. At the top, there are four numbered steps: 1. Basics (selected), 2. Advanced, 3. Resiliency Manager Network, and 4. Infrastructure Management Network. Below the steps, the 'Project details' section asks to select a subscription and resource group. The 'Subscription' dropdown is set to 'VRP-team Subscription' and the 'Resource group' dropdown is set to '(New) VRP', with a 'Create new' link below it. The 'Instance details' section has the 'Region' dropdown set to 'East US'. The 'Resiliency Platform Bootstrap Inputs' section has the 'Select deployment type' dropdown set to 'All Appliances (fresh install)'. A grey information banner at the bottom states: 'It will deploy a fresh Resiliency Manager, an Infrastructure Management Server (IMS) and a Replication Gateway'. At the very bottom, there are two buttons: '< Previous' and 'Next'.

Table 1-17 Basic settings for Azure deployment

Input field	Description
Select Deployment Type	Select the appropriate deployment type from the given options. By default, all the appliances are selected which will deploy a Resiliency Manager, an IMS, a Replication Gateway, Resiliency Manager upgrade and IMS upgrade appliance. You can select other option as per the requirement.
Password for admin user	Set the password for admin user. The admin user and password is later used for configuring the appliances.
Confirm password	Provide same password for confirmation.
NTP Servers	Resiliency Platform requires NTP servers for synchronizing time across all the appliances. Specify FQDN or IP addresses of one or more (space separated) NTP servers. It is recommended to use 3 or more (an odd number) NTP servers.
Timezone	Timezone to be set for all the Resiliency Platform appliances.
Subscription	Select the subscription of Azure account, to be used for deploying the virtual appliances.
Resource group	Specify the name of an existing resource group or a new resource group to be created.
Location	Select the location where you want to create the appliances.

Figure 1-5 Advanced settings for Azure deployment

The screenshot shows the 'Advanced' settings page in the Azure portal. At the top, there are navigation tabs: 'Basics' (selected), 'Advanced', 'Resiliency Manager Network', 'Infrastructure Management Network', 'Replication Gateway Network Settings', and 'Review + create'. The 'Advanced' tab is active, showing three main sections of settings:

- Resiliency Manager Settings:**
 - Resiliency Manager Instance Name: vrp-8M
 - Resiliency Manager Hostname: vrp-rm
 - Resiliency Manager Instance Size: 1x Standard D8s v3 (8 vcpus, 32 GB memory)
 - Data Disk Size in GB: 100
 - Data disk type: Premium SSD
- Infrastructure Management Server (IMS) Settings:**
 - IMS Instance Name: vrp-ims
 - IMS Hostname: vrp-ims
 - IMS Instance Size: 1x Standard F8s (8 vcpus, 16 GB memory)
 - Data Disk Size in GB: 40
 - Data disk type: Premium SSD
- Replication Gateway Settings:**
 - Replication Gateway Instance Name: vrp-gw1
 - Replication Gateway Hostname: vrp-gw1
 - Replication Gateway Instance Size: 1x Standard F8s (8 vcpus, 16 GB memory)
 - Staging Disk Size in GB: 50
 - Staging disk type: Premium SSD
 - Enable Ultra disk compatibility:
 - Availability Zone: (empty)

Table 1-18 Advanced settings for Azure deployment

Input field	Description
Resiliency Manager Settings	
Instance Name	Name of the appliance instance
Host name	Hostname of the appliance. This is changed and allocated from DHCP details
Instance Size	Size of the appliance instance
Data Disk Size in GB	Additional disk of minimum 100 GB required for Resiliency Manager
Infrastructure Management Server (IMS) Settings	
Instance Name	Name of the appliance instance

Table 1-18 Advanced settings for Azure deployment (*continued*)

Input field	Description
Host name	Hostname of the appliance. This is changed and allocated from DHCP details
Instance Size	Size of the appliance instance
Data Disk Size in GB	Additional disk of minimum 40 GB required for IMS
Replication gateway Settings	
Instance Name	Name of the appliance instance
Host name	Hostname of the appliance. This is changed and allocated from DHCP details
Instance Size	Size of the appliance instance
Staging disk size in GB	Additional disk of minimum 50 GB required for Replication Gateway
Staging disk type	Choose staging disk type for Replication Gateway. Learn more about the managed disk types. Managed disk types
Enable Ultra disk compatibility	Choose this check box if you want to support ultra disks on your Replication Gateway instance
Availability Zone	Specify an availability zone if the selected region for deploying the Replication Gateway has availability zones. Keep this field empty if the region does not support availability zones. The deployment can fail if an incorrect input is provided. See Using Azure ultra disks

Figure 1-6 Network settings for Azure deployment

Basics
 Advanced
 Resiliency Manager Network
 Infrastructure Management Network

Is Multi-NIC Deployment? Yes
 No

i Select virtual network where you want to deploy Resiliency Manager. We do not support creation of virtual network as part of deployment. Make sure you select one of the existing networks.

Configure virtual networks

Network Name *
[Create new](#)

Subnet *

Is the NIC eth0 behind NAT? Yes
 No

Table 1-19 Network settings for Azure deployment

Input field	Description
Is Multi-NIC Deployment?	You can deploy Resiliency Manager, IMS, or Replication Gateway with multiple network interfaces. These interfaces can be used for different communication purpose. Select Yes if required..
Are Both NICs in Same Subnet	Select Yes if multiple interfaces are part of same subnet of a virtual network.
Network Name	Specify the name of an existing network.
Select Subnet	Select a subnet to be associated with the virtual appliance.

Table 1-19 Network settings for Azure deployment (*continued*)

Input field	Description
Role of NICs'	In case of multiple interfaces, select NIC for its communication purpose. Make sure you select 'eth0' for at least one intent. If it is single NIC deployment then same NIC is used.
Is the NIC eth0 behind NAT?	Select Yes if NAT is used in your network setup and provide NAT IP and fully qualified NAT hostname.

Similar input is required for configuring IMS, Replication Gateway.

Deploy virtual appliances in Azure Stack using Azure Stack Marketplace

To know about virtual appliance deployment in Veritas Resiliency Platform

A few constraints are applied when you deploy Resiliency Platform in Azure Stack:

See [“Constraints when you deploy the virtual appliances in Azure Stack”](#) on page 405.

Note that, Azure Stack is supported with Resiliency Platform version 3.5 and above.

Below are few prerequisites before deploying in Azure Stack Marketplace:

- Your Azure Stack deployment must have internet connectivity.
- Azure Stack needs to be registered with Azure.

To download Azure Marketplace items to Azure Stack:

- 1 Sign in to the Azure Stack administrator portal.
- 2 Review the available storage space before downloading marketplace items. Later, when you select items for download, you can compare the download size to your available storage capacity.

To review available space, navigate to **Region management** and select the region you want to explore. Later, go to **Resource Providers > Storage**.

- 3 Open Azure Stack Marketplace and connect to Azure. To do so, navigate to: **Marketplace management service** and select Marketplace items, and then select **Add from Azure**.
- 4 Select the **Veritas™ Resiliency Platform Express Install** from the list and then select **Download**. The download time may vary depending upon the internet connectivity. After the download completes, you can deploy the new marketplace item either as an Azure Stack operator or a user.

After the download is complete, to deploy the virtual appliances, refer to the below topic:

See [“Deploying the virtual appliances in Azure through Azure Marketplace”](#) on page 405.

Deploying the virtual appliances in vCloud

To know about virtual appliance deployment in Veritas Resiliency Platform:

You need to deploy at least three Resiliency Platform virtual appliances in your data center in vCloud and then configure Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway.

To deploy the Resiliency Platform virtual appliances in vCloud

1 Recommendations:

While deploying the virtual appliances, ensure the following:

- Deploy the Resiliency Manager and IMS together in a vApp. This vApp will function as a Main Management vApp.
- Deploy the Replication Gateway in a vApp other than the one in which you have deployed the Resiliency Manager and IMS.
- Deploy the Replication Gateway in the Organization where you want to migrate the virtual machines from your on-premises data center.

2 Open vCloud Director, log into the Organization where you want to deploy the Resiliency Manager, IMS, or Replication Gateway.

3 Click **My Cloud**, in the left pane click **vApps**.

4 Use one of the following methods to deploy the OVA in vCloud:

- **Add vApp From Catalog:** Select the OVA from the catalog while adding the vApp.
- **Build a new vApp:** This method can be used to deploy multiple virtual machines at a time. Select the OVA from the catalog while building a new vApp.

Follow the documentation of VMware to deploy the OVA in vCloud using any one of the above-mentioned methods.

Ensure that the details on **Customize Hardware** page match the system resource requirements mentioned:

- 5 After the successful deployment of the OVA, configure the appliance as a Resiliency Manager, IMS, or Replication Gateway.

See “[Configuring the Resiliency Manager or IMS](#)” on page 439.

See “[Configuring the Replication Gateways](#)” on page 445.

Prerequisites for deploying the virtual appliances in vCloud Director

Following are the deployment prerequisites for recovery to vCloud:

- Configure vCloud to have unique BIOS UUID for all the virtual machines when instantiating from a vApp template. By default, all the virtual machines that are created when you deploy a vApp template are assigned the same BIOS UUID. To change this default behavior, follow the steps given in the VMware knowledge bank article:
<https://kb.vmware.com/kb/2002506>
- Set the value of the **disk.enableUUID** configuration parameter as **True** for the following templates uploaded to catalog:
 - Veritas_Resiliency_Platform_VMware_vCloud_vApp_Template_Site1
 - Veritas_Resiliency_Platform_VMware_vCloud_vApp_Template_Site2
 - Veritas_Resiliency_Platform_VMware_vCloud_vApp_Template_Win_Site1
 - Veritas_Resiliency_Platform_VMware_vCloud_vApp_Template_Win_Site2
 - Resiliency Platform Data Mover
- If you want to get the virtual appliances deployed on some specific datastores in vCloud, then you need to create a storage policy. If you do not create a storage policy, the virtual appliances are deployed on any of the available datastores.

Uploading OVA files to catalog

The cloud administrator needs to create a catalog and upload the OVA files in to the catalog.

To upload the OVA files to catalog

- 1 Log into the vCloud service provider Organization in vCloud Director.
- 2 Create a catalog and share the catalog with the tenant Organizations.
- 3 Upload the OVA files of all the Resiliency Platform appliances and all the vApp templates in to the catalog that you have created.
- 4 Ensure that the template vApp names are set to the following:

- VRP_VAPP_TEMPLATE_SITE1
 - VRP_VAPP_TEMPLATE_SITE2
 - VRP_VAPP_TEMPLATE_WIN_SITE1
 - VRP_VAPP_TEMPLATE_WIN_SITE2
- 5 Make the OVAs available to the tenants for deployment of the virtual appliances.
 - 6 Edit the settings for Data Mover and VRP_VAPP_TEMPLATE virtual machines from underlying vCenter servers to set the **disk.enableUUID** parameter as **True**.

Deploying the virtual appliances in Orange Recovery Engine

To know about virtual appliance deployment in Veritas Resiliency Platform:

Following is an overview of the key steps that are performed for deploying the Resiliency Platform virtual appliances in Orange Recovery Engine. You need to perform these steps for each of the virtual appliances:

Table 1-20 Overview of deployment process in Orange Recovery Engine

Step	Action	Description
1	Ensure that the prerequisites for deploying the virtual appliances in Orange Recovery Engine are met.	See “Prerequisites for deploying the virtual appliances in Orange Recovery Engine” on page 416.
2	Download the zip files required for deploying the Resiliency Platform virtual appliance in Orange Recovery Engine.	
3	Upload the image files to Object Storage Service (OBS) using OBS Browser tool. Alternatively, you can use any S3 compliant tool to perform this task.	See “Uploading image files to OBS” on page 416.
4	Using the Orange Recovery Engine console, register the system disk image as a private image.	See “Registering the disk image as a private image” on page 417.
5	Using the Orange Recovery Engine console, launch the ECS instance and attach the data disk to the system disk image that you registered in the previous step.	See “Launching the system disk image instance” on page 418.
6	Restart the ECS.	

Prerequisites for deploying the virtual appliances in Orange Recovery Engine

Following are the prerequisites for deploying the Resiliency Platform virtual appliances in Orange Recovery Engine:

- Ensure to create virtual private network (VPC) and subnets in Orange Recovery Engine.
- Ensure to configure an external DNS server in Orange Recovery Engine. This DNS server is used while configuring the virtual appliances. The host names of IMS and Replication Gateway at source data center should be resolvable from the target data center, and hostname of Resiliency Manager should be resolvable from the source data center.
- Follow the documentation of Orange Recovery Engine to create the required security groups. Make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. You may choose to restrict incoming internet access on these virtual appliances.
- OBS Browser tool needs to be installed on the system through which you are planning to deploy the virtual appliances in Orange Recovery Engine. Alternatively, you can use any S3 compliant tool to perform this task.
- The user who is going to deploy the virtual appliances must have read and write access control list (ACL) on the OBS bucket that is used for deploying the virtual appliances.
- Ensure that the bucket policy settings for the bucket that is used for deploying the virtual appliances do not restrict uploading files to the bucket.

See [“Deploying the virtual appliances in Orange Recovery Engine”](#) on page 415.

Uploading image files to OBS

Before you start deploying the Resiliency Platform virtual appliances in Orange Recovery Engine, you need to upload the downloaded image files to Object Storage Service (OBS).

To upload image files to OBS

- 1 Log in to OBS Browser.
- 2 Select a bucket with **Standard** storage class. If such a bucket does not exist, create a bucket.
- 3 Select the bucket where you want to upload the files.

- 4 Click **Upload** and then click **Upload File**.
- 5 In the next wizard, select the image files to be uploaded.
- 6 Ensure that **Standard** storage class is selected.
- 7 Click **OK** to upload the files. If an upload operation gets suspended or fails, restart the operation. The task will be resumed from the point where it got suspended last time.

See [“Deploying the virtual appliances in Orange Recovery Engine”](#) on page 415.

Registering the disk image as a private image

As a part of deploying the Resiliency Platform virtual appliances in Orange Recovery Engine, you need to register the system image file as a private image.

To register the system image file as a private image

- 1 Log in to the Orange Flexible Engine management console and go to **Image Management Service**.
- 2 On the **Image Management Service** page, click **Create Private Images** tab and click **Create Image**.
- 3 On the **Create Image** page, under **Image Type and Source**, select **System Disk Image** as type.
- 4 For **Source**, select **Image file** to use an external image file.
- 5 Select the bucket where you have uploaded the image files. Navigate to the system image file that you want to register as a private image and select it:
 - System disk for Resiliency Manager:
`Resiliency_Platform_RM_v36-disk1.qcow2`
 - System disk for IMS: `Resiliency_Platform_v36-disk.qcow2`
 - System disk for Replication Gateway:
`Resiliency_Platform_DM_v36-disk1.qcow2`
- 6 Under **Image Information**, ensure that **ECS system disk image** is selected as **Function**.
- 7 For **System Disk (GB)**, enter a minimum value of 40 GB.
- 8 Enter a name for the image.
- 9 Click **Create Now** and review the information displayed on the next page.
- 10 Click **Submit** button to agree to the Orange Flexible Engine Image Disclaimer and click **Submit**.

See [“Deploying the virtual appliances in Orange Recovery Engine”](#) on page 415.

Launching the system disk image instance

After registering the system disk as a private image in Orange Recovery Engine, you need to launch the system disk image instance and attach the data disk to the system disk. For Resiliency Manager and IMS, you need to attach the data disk image that you had downloaded earlier. For Replication Gateway, you attach an extra disk of minimum 50 GB.

After registering the system disk as a private image in Orange Recovery Engine, you need to launch the system disk image instance and attach the data disk to the system disk. For Resiliency Manager, IMS and Replication Gateway attach an extra external disk with below minimum size:

Resiliency Manager: 100 GB

IMS: 40 GB

Replication Gateway: 50 GB

To launch the system disk image instance

- 1 In the Orange Recovery Engine management console, go to **Image Management Service**.
- 2 Select the private image that you had registered in the earlier step.
- 3 Click **Apply for Server** displayed next to the name of the private image.
- 4 On the next page, select the billing mode that suits your requirements.
- 5 Ensure that the desired region and Availability Zone are selected.
- 6 Specify the type and select a flavor that matches the system resource requirements for various Resiliency Platform components.

Note: Select a flavor which is KVM hypervisor based for Replication Gateway, as based hypervisor is not supported.

For more information on hypervisor flavor,

see https://docs.prod.cloud.oracle.com/en-us/usermanual/ecs/en-us_topic_0035470101.html

- 7 In the **Disk** section, select **High I/O** or **Ultra-high I/O** for the system disk. It is strongly recommended to select **Ultra-high I/O** for the disk.
- 8 Ensure that the required image name is selected in the drop down list for image name.
- 9 Click **Add data Disk** and then select **High I/O** or **Ultra-high I/O** for the data disk.

- 10 You can attach external disk of below minimum size to the virtual appliances:
 - Resiliency Manager: 100 GB
 - IMS: 40 GB
 - Replication Gateway: 50 GB
 - 11 Select the **VPC** that you want to use for the deployment.
 - 12 Enter the value for **Security Group** and **NIC**. It is strongly recommended not to assign any public IP or **EIP** to any of the Resiliency Platform appliances.
 - 13 You need to create a key pair or select existing key pair if you have already created it.
 - 14 Enter a name for the ECS. Click **Next**.
 - 15 Verify the details displayed under **Configuration** and click **Create Now**.
- See [“Deploying the virtual appliances in Orange Recovery Engine”](#) on page 415.

Deploying the virtual appliance through VMware vSphere Client

To know about virtual appliance deployment in Veritas Resiliency Platform:

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

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

11 Power on the virtual machine.

To deploy Resiliency Platform through VMware vSphere Web Client

- 1** Download the VMware supported OVA file for the Resiliency Platform virtual appliance on a system where VMware vSphere Web Client is installed.
- 2** In the VMware vSphere Web Client, click **vCenter Servers** and select a vCenter Server. Click **Actions > Deploy OVF template**.
- 3** Select the source location of the Resiliency Platform virtual appliance OVA file.
- 4** Specify a name and location for the deployed template.
- 5** Select a cluster, host, vApp, or resource pool in which to run the deployed template.
- 6** Select a location to store the files for the deployed template.
- 7** Configure the networks the deployed template should use.
- 8** Review the virtual machine configuration and click **Finish**.
- 9** If you want to use DHCP as your network, enter the MAC address of the appliance in the DHCP server. For information on obtaining the MAC address of the appliance, see the documentation of VMware vSphere client.
- 10** Power on the virtual machine.

You can now configure the Resiliency Platform component.

Note: When you deploy the Resiliency Platform Data Mover appliance from vCenter Server 6.5, a warning related to advanced configuration may be displayed. You can ignore the warning and click next to accept the advanced configuration options.

For more information on VMware vSphere Desktop Client or VMware vSphere Web Client, refer to VMware documentation.

See [“About configuring the Resiliency Platform components”](#) on page 437.

Deploying the virtual appliance through Hyper-V Manager

To know about virtual appliance deployment in Veritas Resiliency Platform:

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.
- 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 for IMS or Replication Gateway and 32 GB RAM for Resiliency Manager.
- 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 the VHD file of the Resiliency Platform virtual appliance as **IDE Controller 0**.
- 11 Click **Apply**, and then click **OK**.
- 12 If you want to use DHCP as your network, enter the MAC address of the appliance in the DHCP server. For information on obtaining the MAC address of the appliance, see the documentation of Hyper-V Manager.
- 13 Right-click the name of the virtual machine and select **Start** to power on the virtual machine.

You can now configure the Resiliency Platform component.

See [“About configuring the Resiliency Platform components”](#) on page 437.

Deploying the virtual appliances in Google Cloud Platform (GCP) through GCP Marketplace

Veritas Resiliency Platform enables you to deploy the virtual appliances in Google Cloud Platform through GCP Marketplace. There are four offerings available for deploying the virtual appliances using the templates:

- **Veritas™ Resiliency Platform Express Install 10.0:** Installs a Resiliency Manager, an IMS, and a Replication Gateway appliance of version 10.0.
- **Veritas™ Resiliency Manager 10.0:** Installs a Resiliency Manager appliance of version 10.0.

- **Veritas™ Infrastructure Management Server 10.0:** Installs an Infrastructure Management Server appliance of version 10.0.
- **Veritas™ Replication Gateway 10.0:** Installs a Replication Gateway appliance of version 10.0.

To deploy the virtual appliances in GCP using the templates

1 Prerequisites:

Ensure that the prerequisites are met. Refer to [API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace](#), for deploying the Resiliency Manager and IMS in Google Cloud Platform Marketplace.

Refer See “[Prerequisites for deploying the virtual appliances in Google Cloud Platform](#)” on page 429.

Refer See “[Ports required for recovery of assets to Google Cloud Platform](#)” on page 430.

- 2 Go to the **GCP Marketplace** and search for the offerings.
- 3 Select the desired offerings from the four available offers. GCP marketplace lets you launch the selected offering. Click **Next**.
- 4 On the next page, provide the values for the input fields: See “[Providing inputs for Resiliency Platform template](#)” on page 423.
- 5 Click **Deploy** to start the deployment of resources. Once the deployment starts you will see list of resources being creating as part of deployment.

Next Steps:

SSH to virtual appliance to set admin password:

Once the deployment succeeds you need to SSH to all the servers and change the default password. Resiliency Platform virtual appliances comes with default password for ‘admin’ user and expires at first login. Follow the below steps:

1. Select the virtual machine and enable port 22 in firewall policy if not selected as part of deployment.
2. Login to the virtual machine using ‘Connect to serial console’ with username ‘admin’ and with default password or SSH public key if you have provided during deployment.

Access Resiliency Manager Web UI:

Once you have set the password you can access the Resiliency Manager Web UI using URL: https://<Resiliency_manager_hostname>/ and perform the operations.

Providing inputs for Resiliency Platform template

You need to provide inputs for creating instances using templates. Some of the fields get auto populated with the default value, you can change the values if required. For rest of the parameters, you need to enter a valid value.

Table 1-21 Input required for creating configuration for Resiliency Manager

Field	Description
Instance Count	Number of Resiliency Managers to be created. Only one instance per deployment is currently supported.
Machine type and Series	Select configuration of the instance. Minimum 8 vCPU and 32 GB RAM is required.
Boot disk size in GB	Shows the OS disk size. For the instance 50 GB is required
Boot Disk type	Select disk type for OS disk
Data disk type	Select disk type of data disk. The instance comes with one data disk.
Data disk size in GB	Disk size for data disk. The default value is 100, you can increase the size as required.
Network interfaces	Select Network configuration for NICs. Resiliency Manager allows to have max 2 NICs
Allow TCP port 22 from the Internet	Select checkbox if you want to enable SSH port on the VM, also provide CIDR block which needs to be allowed by default. If kept empty, then SSH will be allowed to all, i.e. 0.0.0.0/0
Instance Name	Name of the instance to be used. If DNS is not configured, then this will be set as short hostname.
Create role with ID VRP_ROLE_RM	Provide custom role with permissions.
Create service account	Provide service account for IAM policy with the role.
Service Account name	Provide service account name having following permissions required for Resiliency Manager to work: <ul style="list-style-type: none"> ■ compute.regions.list ■ storage.buckets.list ■ runtimeconfig.configs.create ■ runtimeconfig.variables.create ■ runtimeconfig.waiters.create

Table 1-21 Input required for creating configuration for Resiliency Manager
(continued)

Field	Description
NIC to be used for communication with other Resiliency Managers	In case of multiple NICs communication with other Resiliency Manager can be restricted on a particular NIC. Select an Interface to be used for the communication.
NIC to be used for communication with IMS	In case of multiple NICs communication with IMS can be restricted on a particular NIC. Select an Interface to be used for the communication.
NIC to be used for accessing the User Interface	In case of multiple NICs Web UI can be restricted on a particular or allowed on all NIC. Select an Interface to be used for the communication.
NIC to be used as default gateway	Select the interface to be used as default gateway for external communication
Is the eth0 Network Interface behind NAT?	If NAT is configured for eth0 interface, then select 'True'. You need to provide NAT hostname and IP address.
Resiliency Manager eth0 NAT Hostname (Optional)	Provide NAT hostname for eth0 if NAT is configured.
Resiliency Manager eth0 NAT IP (Optional)	Provide NAT IP for eth0 if NAT is configured.
Is the eth1 Network Interface behind NAT?	If NAT is configured for eth1 interface, then select 'True'. You need to provide NAT hostname and IP address.
Resiliency Manager eth1 NAT Hostname (Optional)	Provide NAT hostname for eth1 if NAT is configured.
Resiliency Manager eth1 NAT IP (Optional)	Provide NAT IP for eth1 if NAT is configured.

Table 1-22 Input required for creating configuration for Infrastructure Management Server (IMS)

Field	Description
Instance Count	Number of IMS to be created. Only one instance per deployment is currently supported.
Machine type and Series	Select configuration of the instance. Minimum 8 vCPU and 32 GB RAM is required.
Boot disk size in GB	Shows the OS disk size. For the instance 30 GB is required
Boot Disk type	Select disk type for OS disk

Table 1-22 Input required for creating configuration for Infrastructure Management Server (IMS) *(continued)*

Field	Description
Data disk type	Select disk type of data disk. The instance comes with one data disk.
Data disk size in GB	Disk size for data disk. Default is 40 but you can increase as required.
Network interfaces	Select Network configuration for NICs. Resiliency Manager allows to have max 2 NICs
Allow TCP port 22	Select checkbox if you want to enable SSH port on the VM, also provide CIDR block which needs to be allowed by default. If kept empty then SSH will be allowed to all, i.e. 0.0.0.0/0
Instance Name	Name of the instance to be used. If DNS is not configured, then this will be set as short hostname.
Create role with ID VRP_ROLE_IMS	Provide custom role with permissions.
Create service account	Provide service account for IAM policy with the role.
Service Account name	Provide service account name having following permissions required for operations. Check the prerequisites list for required permissions.
NIC to be used for communication with Resiliency Managers	In case of multiple NICs communication with other Resiliency Manager can be restricted on a particular NIC. Select an Interface to be used for the communication.
NIC to be used for communication with Replication Gateway	In case of multiple NICs communication with Replication Gateway can be restricted on a particular NIC. Select an Interface to be used for the communication.
Network Interface to be used as the default gateway	Select the interface to be used as default gateway for external communication
Is selected NIC to communicate Resiliency Manager behind NAT?	NAT can be configured on the NIC which is used for communication with Resiliency Manager. Select True if NAT is configured on the NIC and provide NAT hostname and IP
NAT Hostname (Optional)	Provide NAT hostname if NAT is configured.
NAT IP (Optional)	Provide NAT IP if NAT is configured.

Table 1-23 Input required for creating configuration for Replication Gateway

Field	Description
Instance Count	Number of IMS to be created. Only one instance per deployment is currently supported.
Machine type and Series	Select configuration of the instance. Minimum 8 vCPU and 32 GB RAM is required.
Boot disk size in GB	Shows the OS disk size. For the instance 30 GB is required
Boot Disk type	Select disk type for OS disk
Data disk type	Select disk type of data disk. The instance comes with one data disk.
Data disk size in GB	Disk size for data disk. Default is 50 but you can increase as required.
Network interfaces	Select Network configuration for NICs. Resiliency Manager allows to have max 2 NICs
Allow TCP port 22	Select checkbox if you want to enable SSH port on the VM, also provide CIDR block which needs to be allowed by default. If kept empty then SSH will be allowed to all, i.e. 0.0.0.0/0
Instance Name	Name of the instance to be used. If DNS is not configured, then this will be set as short hostname.
Create role with ID VRP_ROLE_GW	Provide custom role with permissions.
Create service account	Provide service account for IAM policy with the role.
Service Account name	Provide service account name having following permissions required for operations. Check the prerequisites list for required permissions.
NIC to be used for communication with peer Replication Gateway	In case of multiple NICs communication with other Resiliency Manager can be restricted on a particular NIC. Select an interface to be used for the communication.
NIC to be used for communication with IMS	In case of multiple NICs communication with Infrastructure Management Server can be restricted on a particular NIC. Select an Interface to be used for the communication.
NIC to be used for communication with workload Virtual Machines	In case of multiple NICs communication with workload can be restricted on a particular NIC. Select an interface to be used for the communication.

Table 1-23 Input required for creating configuration for Replication Gateway
(continued)

Field	Description
NIC to be used as the default gateway	Select the interface to be used as default gateway for external communication
NIC to communicate with peer Replication Gateway behind NAT?	NAT can be configured on the NIC which is used for communication with peer Replication Gateway. Select True if NAT is configured on the NIC and provide NAT hostname and IP address.
NAT Hostname (Optional)	Provide NAT hostname if NAT is configured.
NAT IP (Optional)	Provide NAT IP address if NAT is configured.

Table 1-24 Common inputs

Fields	Inputs required
SSH public key to this instance	Provide public SSH key to allow key based authentication to the instances. The same key will be configured on all the instances
NTP Server	Provide comma separated one or more NTP servers' hostname or IP address. All the Resiliency Platform virtual appliances must in time sync to perform operations.
Timezone	Select timezone for the instances.

API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace.

Following are the API permissions required to deploy Resiliency Manager and IMS using Google Cloud Platform through Marketplace. Before referring to the below table, you need to know there are some more permissions. Refer to the topic [API permissions required for Google Cloud Platform](#)

Table 1-25 Permissions for deploying Resiliency Manager and IMS using Google Cloud Platform through Marketplace.

Service name	Permissions
runtimeconfig	runtimeconfig.configs.create
	runtimeconfig.variables.create
	runtimeconfig.waiters.create

Prerequisites for deploying the virtual appliances in Google Cloud Platform

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

1. Follow the documentation of Google Cloud Platform to create the required network tags. Make sure that the network tags meet the network and port requirements mentioned in the Resiliency Platform documentation are open for communication. If you deploy Resiliency Platform components through Google Cloud Platform marketplace, the required network tags are automatically created.

See [“Ports required for recovery of assets to Google Cloud Platform”](#) on page 430.
2. Create individual Service Accounts for Resiliency Manager and IMS with certain permissions. These service accounts are used for authenticating the operations performed by the Resiliency Platform components in Google Cloud Platform.

See [“API permissions required for Google Cloud Platform”](#) on page 460.
3. If you deploy Resiliency Platform through Google Cloud Platform marketplace, then the required service accounts are automatically created through Google Cloud Platform template used by marketplace deployment.
4. Ensure that there is direct communication between the premise network and the Google Cloud Platform network. It is recommended to use VPN for Google Cloud Platform environment.
5. Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.

Deploying the virtual appliances in Google Cloud Platform using OVA files

This topic explains about the key steps that are performed for deploying the Resiliency Platform virtual appliances in Google Cloud Platform (GCP). To know about virtual appliance deployment in Veritas Resiliency Platform, refer

Table 1-26 Overview of deployment process in GCP

Step	Action	Description
1	Ensure that the prerequisites for deploying virtual appliances in GCP are met.	Refer See " Prerequisites for deploying the virtual appliances in Google Cloud Platform " on page 429.
2	Upload the OVA files to Google Cloud Storage.	Refer See " Uploading the OVA file using web-based method " on page 433. Refer See " Uploading the OVA file using command-line method " on page 434.
3	Create image from the uploaded OVA file.	Refer See " Creating Image using web-based method " on page 435. Refer See " Creating Image using command-based method " on page 435.
4	Launch the instances of virtual appliances to deploy Resiliency Manager, Infrastructure Manager (IMS), and Replication Gateway	Refer See " Launching the instances of virtual appliances " on page 436.

Prerequisites for deploying the virtual appliances in Google Cloud Platform

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

1. Follow the documentation of Google Cloud Platform to create the required network tags. Make sure that the network tags meet the network and port requirements mentioned in the Resiliency Platform documentation are open for communication. If you deploy Resiliency Platform components through Google Cloud Platform marketplace, the required network tags are automatically created.

See "[Ports required for recovery of assets to Google Cloud Platform](#)" on page 430.

2. Create individual Service Accounts for Resiliency Manager and IMS with certain permissions. These service accounts are used for authenticating the operations performed by the Resiliency Platform components in Google Cloud Platform.

- See [“API permissions required for Google Cloud Platform”](#) on page 460.
3. If you deploy Resiliency Platform through Google Cloud Platform marketplace, then the required service accounts are automatically created through Google Cloud Platform template used by marketplace deployment.
 4. Ensure that there is direct communication between the premise network and the Google Cloud Platform network. It is recommended to use VPN for Google Cloud Platform environment.
 5. Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.

Ports required for recovery of assets to Google Cloud Platform

Following is the list of ports required for recovery of assets to Google Cloud Platform:

Table 1-27 Ports required for recovery of assets to Google Cloud Platform

Ports for recovery to Google Cloud Platform
See “Ports required for Resiliency Manager” on page 78.
See “Ports required for IMS” on page 80.
See “Ports required for Replication Gateway used for recovery to cloud data center” on page 81.
See “Ports required for hosts” on page 82.

Ports required for Resiliency Manager

Table 1-28 Resiliency Manager

Ports used	Purpose	For communication between	Direction	Protocol
22	Used for Remote access to Resiliency Manager Shell Console.	Resiliency Manager and the hosts	In-bound	TCP
25	Used for SMTP mail server.	Resiliency Manager and SMTP server	Out-bound	TCP
53	Used for DNS name resolution.	DNS server and Resiliency Manager	Out-bound	TCP, UDP
123	Used for NTP synchronization.	Resiliency Manager and the NTP server	Out-bound	UDP

Table 1-28 Resiliency Manager (*continued*)

Ports used	Purpose	For communication between	Direction	Protocol
162	Used for configuring SNMP port.	Clients listening to the traps sent by Resiliency Manager	Out-bound	UDP
389	Used for communication with LDAP/AD server.	Resiliency Manager and LDAP/AD server	Out-bound	TCP, user provided
443	Used for SSL communication.	Resiliency Manager and web browser	In-bound	HTTPS, TLS 1.2
636	Used for communication with LDAP/AD server.	Resiliency Manager and LDAP/AD server	Out-bound	TCP with SSL / TLS, user provided
7000 and 7001	Used for database replication.	Between Resiliency Managers in case of multiple Resiliency Managers	Bi-directional	TCP with SSL / TLS1.1+
8000	Used only if you want to use Object Storage for replication to AWS. Or in case of multiple Resiliency Manager setup upgrade.	Resiliency Manager and lambda functions, or among multiple Resiliency Managers while upgrading	Bi-directional	TCP
14161	Used for adding IMS by providing username and password.	Resiliency Manager and IMS	Out-bound	HTTPS, TLS 1.2
14176	Used for communication between the Resiliency Manager and IMS.	Resiliency Manager and IMS	In-bound	HTTPS, TLS 1.2
14176	Used for communication between Resiliency Managers in case of multiple Resiliency Managers.	Between Resiliency Managers in case of multiple Resiliency Managers	Bi-directional	HTTPS, TLS 1.2

If you do not have multiple Resiliency Manager setups and only want to use Object Storage for replication to AWS, then port 8000 need not be opened on the VPN.

Ports required for IMS

Table 1-29 Infrastructure Management Server (IMS)

Ports used	Description	For communication between	Direction	Protocol
22	Used for Remote access to IMS Shell Console	Client terminals to IMS	In-bound	TCP
22	Used for remote deployment of the packages on remote Linux host from IMS	IMS to hosts to be protected	Out-bound	TCP
53	Used for DNS name resolution	DNS server and IMS in the data center	Out-bound	TCP, UDP
88	Used for authenticating DNS server with Kerberos keys	DNS server and IMS	Out-bound	TCP, UDP
123	Used for NTP synchronization	IMS and the NTP server	Out-bound	UDP
443	Used for communication with vCenter using vSphere SDK This is the default value for the port that will be used for accessing vcenter for discovery and operations. A non-default port must be configured on the vcenter first before it can be configured in IMS.	IMS and VMware vCenter server	Out-bound	HTTPS, TCP
750	Used for authenticating DNS server with Kerberos keys	DNS server and IMS	Out-bound	TCP, UDP

Table 1-29 Infrastructure Management Server (IMS) *(continued)*

Ports used	Description	For communication between	Direction	Protocol
5634	Used for communication between IMS and the hosts that are added to Resiliency Platform	IMS and the hosts	Bi-directional	HTTPS, TLS 1.2
14161	Used for adding IMS by providing username and password	Resiliency Manager and IMS	In-bound	HTTPS, TLS 1.2
14176	Used for communication between the Resiliency Manager and IMS	Resiliency Manager and IMS	Out-bound	HTTPS, TLS 1.2

Ports required for Replication Gateway used for recovery to cloud data center
Table 1-30 Replication Gateway with in-guest replication

Ports used	Description	For communication between	Direction	Protocol
22	Used for remote access to Replication Gateway Shell Console	Replication Gateway and the hosts	In-bound	TCP
53	Used for DNS name resolution	DNS server and Replication Gateway in the data center	Out-bound	TCP, UDP
123	Used for NTP synchronization	Replication Gateway and the NTP server	Out-bound	UDP
443	Used for replication Note: To change port 443 as unidirectional on an on-premise site, enable the outbound traffic on the PORT. This is applicable for Premise to Cloud replication only.	Between the Replication Gateways	Bi-directional	HTTPS, TLS 1.2

Table 1-30 Replication Gateway with in-guest replication (*continued*)

Ports used	Description	For communication between	Direction	Protocol
5634	Used for communication with IMS	IMS and Replication Gateway	Bi-directional	HTTPS, TLS 1.2
8089	Used for replication	Between the Replication Gateways	Bi-directional	TCP
33056	Used for replication	Virtual machine and Replication Gateway	In-bound	TCP

Ports required for hosts

Table 1-31 Ports required for Linux hosts to be protected

Ports used	Description	For communication between	Direction	Protocol
22	Used for remote deployment of the packages on remote Linux host from IMS	IMS and the hosts	In-bound	TCP
53	Used for DNS name resolution	DNS server and host	Out-bound	TCP, UDP
123	Used for NTP synchronization	Host and the NTP server	Out-bound	UDP
5634	Used for communication with IMS	IMS and the hosts	Bi-directional	HTTPS, TLS 1.2

Table 1-31 Ports required for Linux hosts to be protected (*continued*)

Ports used	Description	For communication between	Direction	Protocol
33056	Used for in-guest replication	Virtual machine and Replication Gateway Host and Replication Gateway (when recovering physical machines to VMware environment)	Out-bound	TCP

Table 1-32 Ports required for Windows hosts to be protected

Firewall Rule / Ports used	Description	For communication between	Direction	Protocol
Core Networking - DNS (UDP-Out) / 53	Used for DNS name resolution	DNS server and host	Out-bound	UDP
123	Used for NTP synchronization	Host and the NTP server	Out-bound	UDP
5634	Used for communication with IMS	IMS and the hosts	Bi-directional	HTTPS, TLS 1.2
COM+ Network Access (DCOM-In) / 135	Used for remote deployment on client computer	Windows hosts added to Resiliency Platform and the IMS	In-bound	TCP
Windows Management Instrumentation (WMI-In)	Used for remote deployment on client computer	Windows hosts added to Resiliency Platform and the IMS	In-bound	TCP
File and Printer Sharing (SMB-In) / 445	Used for remote deployment on client computer	Windows hosts added to Resiliency Platform and the IMS	In-bound	TCP

Table 1-32 Ports required for Windows hosts to be protected (*continued*)

Firewall Rule / Ports used	Description	For communication between	Direction	Protocol
33056	Used for in-guest replication	Replication Gateway and virtual machine Replication Gateway and physical machine	Out-bound	TCP

Table 1-33 Ports required for add hyper-v operation

Firewall Rule / Ports used	Description	For communication between	Direction	Protocol
445 (SMB)	Used to add hyper-v operation	Hyper-V server which is hosting hosts to be protected and the IMS	In-bound	TCP

API permissions required for Google Cloud Platform

Following are the permissions required for the roles that you need to create for Resiliency Manager and IMS for recovery to Google Cloud Platform data center. If you are deploying your appliances through Marketplace, then refer to

[API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace.](#)

Table 1-34 API Permissions required for role for Resiliency Manager

Service name	Permissions
compute	compute.regions.list
storage	storage.buckets.list

Table 1-35 API Permissions required for role for IMS

Service name	Permission
cloudkms	cloudkms.cryptoKeyVersions.list
	cloudkms.cryptoKeys.list
	cloudkms.keyRings.list

Table 1-35 API Permissions required for role for IMS (*continued*)

Service name	Permission
compute	compute.addresses.list
	compute.diskTypes.list
	compute.disks.create
	compute.disks.createSnapshot
	compute.disks.delete
	compute.disks.list
	compute.disks.use
	compute.disks.get
	compute.firewalls.list
	compute.globalOperations.list
	compute.images.create
	compute.images.get
	compute.images.useReadOnly
	compute.instances.attachDisk
	compute.instances.create
	compute.instances.delete
	compute.instances.detachDisk
	compute.instances.get
	compute.instances.list
	compute.instances.setMetadata
	compute.instances.setTags
	compute.instances.start
	compute.instances.stop
	compute.machineTypes.list
compute.networks.list	

Table 1-35 API Permissions required for role for IMS (*continued*)

Service name	Permission
	compute.projects.get
	compute.regionOperations.list
	compute.snapshots.create
	compute.snapshots.delete
	compute.snapshots.list
	compute.snapshots.useReadOnly
	compute.subnetworks.list
	compute.subnetworks.use
	compute.zoneOperations.list
	compute.zones.list
	compute.addresses.useInternal
storage	storage.objects.create
	storage.objects.get

Table 1-36 Permissions required to discover Shared VPC for a role for IMS on the host project.

Service name	Permission
compute	compute.networks.list
	compute.firewalls.list
	compute.addresses.list

Note: To share subnets from the host project with configured project, you need to add the Service Account associated with IMS as **Principal** and provide a **Compute Network User role** on the shared subnets.

Uploading the OVA file using web-based method

You can create a Google Cloud Platform cloud storage bucket and upload the ova file to that bucket using a web-based method.

To upload the OVA file using web-based method

- 1 Log into the Google Cloud Platform console.
- 2 Navigate to **Cloud Storage** and click **Create Bucket**.
- 3 Enter a name for the bucket and select the appropriate region.
- 4 Select the Standard storage class.
- 5 Click **Create**.
- 6 Once the bucket is created, open the bucket and click **Upload Files**.
- 7 Select the OVA files from your local disk and it will start uploading the file.

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Uploading the OVA file using command-line method

You need to first create a Cloud Storage bucket in Google Cloud Platform and then upload your ova file to that bucket.

To upload the OVA file using command-line method

- 1 Download and install the [GCP Command Line Interface](#).
- 2 Use the `gsutil mb` command to create a new bucket. Bucket names must be unique and should be DNS compliant:

```
gsutil mb gs://<bucket_name>
```

Where:

- `bucket_name` is the name you want to give your bucket.
- If the request is successful, the command returns the following message:

```
Creating gs://<bucket_name>/...
```

- 3 Upload the OVA file by executing the following command:

```
gsutil cp <object_location> gs://<destination_bucket_name>/
```

Where:

- `object_location` is the local path to your OVA file. For example, `Desktop/RM.ova`.
- `destination_bucket_name` is the name of the bucket to which you are uploading your object. For example, `my-bucket`.

If the request is successful, the response looks like the following example:

```
Operation completed over 1 objects/58.8 KiB
```

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Creating Image using web-based method

Once you upload the OVA files to Google Cloud Storage bucket, you need to create a Image from the OVA files that you have uploaded. This image can be later used to launch the instances for deploying Resiliency Manger, Infrastructure Manager in Google Cloud Platform.

To create image

- 1 Login to the Google Cloud Platform console.
- 2 Navigate to **Images** and click on **Create Image**.
- 3 Enter a name for image.
- 4 Select the **Source** as **Virtual Disk**.
- 5 In Cloud Storage file, browse the OVA uploaded into your bucket.
- 6 In Operating System on virtual disk, select **No operating system. Data only**.
- 7 Click **Create**.

It will start importing image from the selected OVA file.

Note: If your default VPC setting for “Subnet creation mode” from the project is set to “Custom subnets” then create image will not work from Google Cloud Platform console. In that case, you need to use command-based method only.

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Creating Image using command-based method

You can also create image using command-based method:

To create an image by command line

- 1 Open command prompt.
- 2 Execute the following command to create the image.

```
gcloud compute images import <image_name> --source-file  
gs://<bucket_name>/<file_name> --no-guest-environment --no-address  
--network projects/<project_name>/global/networks/default --subnet  
projects/<project_name>/regions/<region>/subnetworks/default
```

Where:

- `bucket_name` is the name of the bucket in which you have uploaded the OVA file.
- `project_name` is name of the project for which you want to create the image.
- `Region` is name of the region for which you need to create the Resiliency Platform virtual appliances.
- `File name` is the name of the ova file which is uploaded.

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Launching the instances of virtual appliances

Once an image gets created, you can use the image to launch instances to deploy the Resiliency Manager and any number of Infrastructure Management Servers (IMS), and Replication Gateways in Google Cloud Platform.

To launch the instances of virtual appliances

- 1 Go to the Google Cloud Platform console and navigate to **Images** under **Storage**.
- 2 Click on the image which you want to select and click **Create Instance**.
- 3 Make sure to select appropriate Machine Type that matches with the system resource requirements mentioned in the documentation.

Network Optimization should be high for the instance.
- 4 Select the required Region and Zone in which you wish to deploy the Resiliency Platform.
- 5 Under **Identity and API access**, select a Service Account created with all the necessary permissions for Resiliency Manager and IMS.

If you deploy the appliances using marketplace, then template deployment will create the required service account and associate it with the instance.
- 6 Under **Networking**, enter the Network Tag created for the particular instance.
- 7 Select the proper Network and Subnetwork in which you want to create the setup.
- 8 Under Disks, you need to attach an extra disk of size 50 GB for Resiliency Manager, 40 GB for IMS and 50 GB for Replication Gateway.
- 9 Click **Create** to launch the instance.

In the **Select an existing key pair** or **Create a new key pair wizard**, you can choose an existing key pair, or create a new one. If you create a new key pair, ensure to click the **Download key pair** button and download. You will need the

private key from this key pair to login as admin user for completing the bootstrap process.

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

About configuring the Resiliency Platform components

After the Veritas Resiliency Platform virtual appliance deployment, you are expected to configure the Resiliency Platform component that you have deployed, through the bootstrap process. The bootstrap process is automatically invoked when you log in to the virtual appliance console for the first time using the admin user login.

Note: There is no sequence required for configuring the Resiliency Platform components. You can configure the components in any sequence on source as well as target data centers. Only after configuring Resiliency Manager, a URL for the Resiliency Platform web console login is provided and then you can access that URL in a web browser to log in to the web console.

The following settings are configured as part of this process to set up the component:

- **Host Network settings:** Settings such as fully qualified hostname (FQDN), IP address, subnet mask, default gateway, and DNS server. Before you use the hostname and the IP address, you need to register them with the DNS server.
- **Appliance settings:** Settings such as NTP server.
- **Product settings:** Configures the virtual appliance as a Resiliency Manager, Infrastructure Management Server (IMS), Replication Gateway.

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

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

Before configuring the component through product bootstrap, ensure that the prerequisites are met.

See [“Prerequisites for configuring Resiliency Platform components”](#) on page 438.

See [“Configuring the Resiliency Manager or IMS”](#) on page 439.

See [“Configuring the Replication Gateways”](#) on page 445.

Prerequisites for configuring Resiliency Platform components

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

- Veritas Resiliency Platform now supports Internet protocol version 6 (IPv6) along with Internet protocol version 4 (IPv4) .
- Before you use the hostname and the IP address in the **Network settings**, you need to register them with the DNS server. Also ensure that the reverse lookup for that IP address works.
- Ensure that the host name corresponding to the IP address is less than 64 characters long. In case of NAT, this is required for host names corresponding to both private and public IP addresses.
- If you plan to use the DHCP server, the DHCP server should be reachable from the subnet and should be able to respond to the subnet where you plan to deploy the product. This requirement is also applicable to static DHCP.
- To use DHCP network, you need to reserve an IPv4 address for the virtual appliance in the DHCP server along with the corresponding MAC address. You cannot configure IPv6 address for the virtual appliance in the DHCP server.
- If you are configuring multiple NICs using DHCP, ensure that same DNS is used to resolve the IP addresses for all NICs.
- In case of multiple Resiliency Managers, you need to either use the same NTP server for configuration or ensure that the NTP servers are properly synchronized.
- Veritas Resiliency Platform supports Linux NTP server. You can also use a public NTP server If there is internet access to the appliance. It is recommended to use a pool (with odd numbers) of time resources for your NTP server. NTP server takes inputs from the available time sources and uses algorithms to find out the correct time. If there are even number of sources and they do not agree, then the algorithm of NTP server fails to make the right decision. Also, it is a better practice to use diversity of reference clocks. You can now configure NTP server using IPv4 or IPv6 address.
- Ensure that the NICs of the virtual appliances have a static MAC address. You can set a static MAC address for the appliance NICs using the virtual machine settings.
- While configuring the virtual appliances on the source data center, ensure that the IP and hostname of Resiliency Manager, IMS, and Replication Gateway can be resolved from the target data center and vice versa.

- In case of a Replication Gateway, make sure to attach an extra disk of at least 50 GB before configuring the Replication Gateway.

See [“Configuring the Resiliency Manager or IMS”](#) on page 439.

Configuring the Resiliency Manager or IMS

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

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

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

To configure the Resiliency Manager or IMS

1 Prerequisites:

See [“Prerequisites for configuring Resiliency Platform components”](#) on page 438.

- #### 2
- In any non-AWS environment, log in to the virtual appliance console or SSH using the following credentials:
 - **Username:** admin
 - **Password:** P@ssw0rd

Note: In Azure environment, the password would be the one provided during deployment.

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

See [“Password policies for Resiliency Platform virtual appliance”](#) on page 451.

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

- In AWS environment, do one of the following:
 - From Linux system:

Use SSH with the private key from the key-pair that you had selected while launching the instance in AWS. For example:

```
ssh -i private_key_file admin@ip_address_of_aws_instance
```

Ensure to modify the permissions for the private key file as 600 before using the file.
 - From Windows system:

Follow the documentation of AWS to connect to the Linux instance from a Windows system using PuTTY.
- 3** Accept the End User License agreement (EULA) to proceed with the configuration.
- 4** In the **Host Network Settings** section, you can configure the appliance network by using DHCP or static IP.

See [“Configuring network settings for Resiliency Manager”](#) on page 441.
See [“Configuring network settings for IMS”](#) on page 443.
- 5** In the **Appliance Settings** section, do the following:
 - Press the Enter key to confirm the use of an NTP server for configuring the date and time.
 - You are required to select the time zone. Follow the instructions as displayed on the virtual appliance console or SSH session to select the correct time zone.
 - Enter the FQDN or IP address of the NTP server. The appliance verifies the NTP server details. If there are any issues, details are displayed on the screen and you are prompted to enter the details again.

You can reset the timezone and NTP server at a later point of time using klish menu. Changing the system settings can affect the product functionality if incorrect values are set.
- 6** In the **Product Settings** section, the virtual appliance is configured as Resiliency Manager or IMS, depending upon the OVA file that you had selected for deployment.

- 7 After a successful product configuration, a message is displayed. If you have configured Resiliency Manager on the virtual appliance, a URL for the Resiliency Platform web console login is provided. You can type the URL in a web browser and log in to the web console.
- 8 If the bootstrap is in AWS environment, you must log in once again using the SSH key and set the admin user password. You need to use this password for subsequent logins to the console.

See [“About configuring the Resiliency Platform components”](#) on page 437.

Configuring network settings for Resiliency Manager

A Resiliency Manager needs to communicate with multiple entities within Veritas Resiliency Platform such as the IMS and another Resiliency Manager in the domain. To facilitate separate communication channels for these communications, Resiliency Platform 10.0 extends support for configuring Resiliency Manager with three Network Interface Cards (NIC).

The Resiliency Manager appliance is shipped with three NICs. You can configure these NICs to be used for the following three communications:

- For communication with Infrastructure Management Server (IMS)
- For communication with other Resiliency Managers and clouds too
- For communication with Product User Interface

If you do not plan to use three separate networks, you do not need to configure three NICs. You can configure one or multiple NICs based on your network layout. So, if you have only one network, you can configure only one NIC and associate all communications to go through the configured NIC.

Since IPv6 network support is provided from Resiliency Platform 3.3.1 version, you can configure the NICs using IPv4 and IPv6 addresses. For more information,

Note: The network configuration of the NICs performed during the bootstrap is final and you cannot edit the network configuration of the NICs at a later point of time.

To configure network settings for Resiliency Manager

- 1 In the **Host Network Settings** section, the bootstrap program first checks for network configuration of all the NICs of the appliance and prints the network details of all the NICs. You are prompted to continue the process of bootstrap. You can see a list of sections that need to be completed as a part of host network settings:
 - Host Network Settings for communication with Infrastructure Management Server

- Host Network Settings for communication to other Resiliency Managers
- Host Network Settings for communication with Product User Interface

2 For each section listed above, do the following:

- Enter the NIC to be used for the communication. The name and MAC address of NIC is displayed.
- The Bootstrap process checks if the NIC is already configured.
 - If the NIC is already configured then the NIC network configuration details are printed and you are prompted to confirm if you want to continue with the printed configuration.
 - If the NIC is not configured or if you do not want to use the existing NIC configuration, then you can choose to use either DHCP protocol or static protocol. In case of static protocol, you need to provide static network details such as IP address, prefix length.

Note: Ensure that appropriate subnet or virtual switch is assigned to the network adapter. Confirm this by matching the MAC address shown in the bootstrap with the one assigned to the virtual machine by the virtualization or cloud technology.

- Confirm if you want to add a static route. You need to set a static route for the interface only if you want the interface to reach a subnet that is different from all the subnets configured on this appliance and the default gateway is unable to communicate with that subnet. In this case, provide the subnet details (in CIDR format) at the input prompt. You can also set a static route to a host. In this case, provide the IP address of the host at the input prompt.
- Confirm if you are in Network Address Translation (NAT) environment and want to configure NAT when the NICs are configured in IPv4 networks only.

Note: Since NAT is not supported for IPv6 address, hence when you configure the virtual appliance using IPv6 address only and both IPv4 and IPv6 address, you are not asked about the NAT configuration. To know more about NAT support in Resiliency Platform

See [“About NAT support in Veritas Resiliency Platform”](#) on page 451.

You can add NAT gateway for communication between multiple Resiliency Managers, only if all the Resiliency Managers are not deployed in the same data center.

- 3 After successful configuration of the first NIC, confirm if you want to use the same NIC for the other communication channel. If you do not want to use the same NIC, perform step 2 for the other two communication channels.

Note: You can select one or more NICs for communication with Product User Interface. If you want to use multiple NICs to access product user interface, enter space separated values of the NICs.

- 4 Enter details for default router and then enter the DNS server details.
- 5 The details of **Host Network Settings** are displayed and you are prompted to confirm. Review the information. If any information is incorrect, choose **n** to go back to the networking inputs page and correct the details. Upon confirmation, the network is configured.

See [“Configuring the Resiliency Manager or IMS”](#) on page 439.

Configuring network settings for IMS

An IMS needs to communicate with multiple entities within Veritas Resiliency Platform such as the Resiliency Manager and the protected hosts. To facilitate separate communication channels for these communications, Resiliency Platform 10.0 extends support for configuring IMS with three Network Interface Cards (NIC).

The IMS appliance is shipped with three NICs. You can configure these NICs to be used for the following two communications:

- For communication with Resiliency Manager
- For communication with the gateways as well as with the hosts to be protected

Remaining one NIC is not configured during the bootstrap. Using the klish menu, you can configure that NIC later for communication between the appliance and any external entity. You can now configure this remaining NIC using IPv6 address with `nic-configuration set` command.

See [“Klish menu options for IMS”](#) on page 598.

If you do not plan to use two separate networks, you do not need to configure the two NICs. You can configure only one NIC or two NICs based on your network layout. So if you have only one network, you can configure only one NIC and associate the two communications to go through the configured NIC.

Since IPv6 network support is provided from Resiliency Platform 3.3.1 version, you can configure the NICs using IPv4 and IPv6 addresses. For more information,

Note: The network configuration of the NICs performed during the bootstrap is final and you cannot edit the network configuration of the NICs at a later point of time.

To configure network settings for IMS

- 1 In the **Host Network Settings** section, the bootstrap program first checks for network configuration of all the NICs of the appliance and prints the network details of all the NICs. You are prompted to continue the process of bootstrap. You can see a list of sections that need to be completed as a part of host network settings:
 - Host Network Settings for communication with Resiliency Manager
 - Host Network Settings for communication with Replication Gateway, Workload Virtual Machines, And Discovery Hosts
- 2 For each section listed above, do the following:
 - Enter the NIC to be used for the communication. The name and MAC address of NIC is displayed.
 - The Bootstrap process checks if the NIC is already configured.
 - If the NIC is already configured then the NIC network configuration details are printed and you are prompted to confirm if you want to continue with the printed configuration.
 - If the NIC is not configured or if you do not want to use the existing NIC configuration, then you can choose to use either DHCP protocol or static protocol. In case of static protocol, you need to provide static network details such as IP address, prefix length.

Note: Ensure that appropriate subnet or virtual switch is assigned to the network adapter. Confirm this by matching the MAC address shown in the bootstrap with the one assigned to the virtual machine by the virtualization technology or cloud technology.

- Confirm if you want to add a static route. You need to set a static route for the interface only. If you want the interface to reach a subnet that is different from all the subnets configured on this appliance and the default gateway is unable to communicate with that subnet. In this case, provide the subnet details (in CIDR format) at the input prompt. You can also set a static route to a host. In this case, provide the IP address of the host at the input prompt.
- Confirm if you are in Network Address Translation (NAT) environment and want to configure NAT when the NICs are configured in IPv4 networks only.

Note: Since NAT is not supported for IPv6 address, hence when you configure the virtual appliance using IPv6 address only and both IPv4 and IPv6 address, you are not asked about the NAT configuration. To know more about NAT support in Resiliency Platform

See [“About NAT support in Veritas Resiliency Platform”](#) on page 451.

You can add NAT gateway for communication between Resiliency Manager and IMS, only if IMS and the Resiliency Manager are not deployed in the same data center.

- 3 After successful configuration of the first NIC, confirm if you want to use the same NIC for the other communication channel. If you do not want to use the same NIC, perform step 2 for the other NIC.
- 4 Enter details for default router and then enter the DNS server details.
- 5 The details of **Host Network Settings** are displayed and you are prompted to confirm. Review the information. If any information is incorrect, choose **n** to go back to the networking inputs page and correct the details. Upon confirmation, the network is configured.

See [“Configuring the Resiliency Manager or IMS”](#) on page 439.

Configuring the Replication Gateways

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

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

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

To configure a Replication Gateway

- 1 Prerequisites:
 - See [“Prerequisites for configuring Resiliency Platform components”](#) on page 438.

- The Replication Gateway configuration requires an extra disk of at least 50 GB, to be used as a staging disk. You can attach this disk before configuration or during the configuration. The default disk size of 50GB lets you protect up to 8 virtual machines and each additional virtual machine requires a disk of 6GB size. You can increase the size of the disk by using `lvm` option of the klish commands:
- 2
- In any non-AWS environment, log in to the virtual appliance console or SSH using the following credentials:
 - **Username:** admin
 - **Password:** P@ssw0rd

Note: In Azure environment, the password would be the one provided during deployment.

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

See [“Password policies for Resiliency Platform virtual appliance”](#) on page 451.

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

- In AWS environment, do one of the following:
 - From a Linux client system:

Use SSH with the private key from the key-pair that you had selected while launching the instance in AWS. For example:

```
ssh -i private_key_file admin@ip_address_of_aws_instance
```

Ensure to modify the permissions for the private key file as 600 before using the file.
 - From a Windows client system:

Follow the documentation of AWS to connect to the Linux instance from a Windows system using PuTTY.
- 3
- Accept the End User License agreement (EULA) to proceed with the configuration.

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

See [“Configuring network settings for Replication Gateway”](#) on page 448.

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

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

- 6** In the **Product Settings** section, the virtual appliance is configured as Replication Gateway.
- 7** You are prompted to confirm if you want to enable FIPS for the Replication Gateway appliance.

See [“About FIPS enablement for Replication Gateway appliance”](#) on page 450.

Note: If you confirm to enable FIPS for the appliance, the appliance will be restarted after finishing the bootstrap process.

- 8** You are prompted to attach an extra disk to the appliance. If you have already attached the extra disk, press **Enter** to confirm. If you have attached more than one extra disk, all disks are listed and you need to select the extra disk that you want to use. If you have not already attached the extra disk, attach the extra disk and then confirm or select the extra disk to be used. Ensure that you attach a thick provisioned disk.

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

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

If the bootstrap is in AWS environment, you must log in once again using the SSH key and set the admin user password. You need to use this password for adding the gateway to the resiliency manager.
- 10 Add the Replication Gateway to an IMS using the Resiliency Manager console.

See [“Adding a Replication Gateway”](#) on page 111.

Configuring network settings for Replication Gateway

A Replication Gateway needs to communicate with multiple entities within Veritas Resiliency Platform such as the IMS, protected hosts, and peer Gateway. To facilitate separate communication channels for all these communications, Resiliency Platform 10.0 extends support for configuring Replication Gateway with multiple Network Interface Cards (NIC).

The Replication Gateway appliance is shipped with four NICs. Out of these four NICs, you can configure three NICs to be used for the following three communications:

- For communication with Infrastructure Management Server (IMS)
- For communication with peer Replication Gateway
- For communication with the virtual machines to be protected

Remaining one NIC is not configured during the bootstrap. Using the klish menu, you can configure that NIC later for communication between the appliance and any external entity. You can now configure this remaining NIC using IPv6 address with `nic-configuration set` command..

See [“Klish menu options for Replication Gateway”](#) on page 612.

If you do not plan to use three separate networks, you do not need to configure all the three NICs. You can configure only one NIC or two NICs based on your network layout. So if you have only one network, you can configure only one NIC and associate all three communications to go through the configured NIC. Likewise, if you have two networks, you can configure two NICs and associate appropriate communications to go through these two NICs.

Since IPv6 network support is provided from Resiliency Platform 3.3.1 version, you can configure the NICs using IPv4 and IPv6 addresses. For more information,

Note: The network configuration of the NICs performed during the bootstrap is final and you cannot edit the network configuration of the NICs at a later point of time.

To configure network settings for Replication Gateway

- 1 In the **Host Network Settings** section, the bootstrap program first checks for network configuration of all the NICs of the appliance and prints the network details of the NICs. You are prompted to continue the process of bootstrap. You can see a list of sections that need to be completed as a part of host network settings:
 - Host Network Settings for communication to Infrastructure Management Server
 - Host Network Settings for communication to Peer Gateways
 - Host Network Settings for communication to virtual machines to be protected
- 2 For each section listed above, do the following:
 - Enter the NIC to be used for the communication. The name and MAC address of NIC is displayed.
 - The Bootstrap process checks if the NIC is already configured.
 - If the NIC is already configured then the NIC network configuration details are printed and you are prompted to confirm if you want to continue with the printed configuration.
 - If the NIC is not configured or if you do not want to use the existing NIC configuration, then you can choose to use either DHCP protocol or static protocol. In case of static protocol, you need to provide static network details such as IP address, prefix length.

Note: Ensure that appropriate subnet or virtual switch is assigned to the network adapter. Confirm this by matching the MAC address shown in the bootstrap with the one assigned to the virtual machine by the virtualization or cloud technology.

- Confirm if you want to add a static route. You need to set a static route for the interface only. If you want the interface to reach a subnet that is different from all the subnets configured on this appliance and the default gateway is unable to communicate with that subnet. In this case, provide the subnet details (in CIDR format) at the input prompt. You can also set a static route to a host. In this case, provide the IP address of the host at the input prompt.
- Confirm if you are in Network Address Translation (NAT) environment and want to configure NAT when the NICs are configured in IPv4 networks only.

Note: Since NAT is not supported for IPv6 address, hence when you configure the virtual appliance using IPv6 address only and both IPv4 and IPv6 address, you are not asked about the NAT configuration. To know more about NAT support in Resiliency Platform

See “[About NAT support in Veritas Resiliency Platform](#)” on page 451.

You can add NAT gateway for communication between the peer gateways.

- 3 After successful configuration of the first NIC, confirm if you want to use the same NIC for other two communication channels. If you do not want to use the same NIC, perform step 2 for rest of the NICs.
- 4 Enter details for default router and then enter the DNS server details.
- 5 The details of **Host Network Settings** are displayed and you are prompted to confirm. Review the information. If any information is incorrect, choose **n** to go back to the networking inputs page and correct the details. Upon confirmation, the network is configured.

See “[Configuring the Replication Gateways](#)” on page 445.

About FIPS enablement for Replication Gateway appliance

Federal Information Processing Standards (FIPS) is a set of standards that describe document processing, encryption algorithms, and other information technology standards for use within the non-military government agencies and by government contractors and vendors who work with these agencies.

Resiliency Platform Data Mover lets you configure encryption of data over WAN (Wide Area Network). It uses the openssl library provided by the operating system vendor RedHat to encrypt the data before transmitting the data. Openssl library shipped with RHEL 6.6 and onwards is certified under the certificate numbers 2446 and 2447:

<http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2015.htm>

Veritas Resiliency Platform provides 128 bit and 256 bit encryption schemes. When the operating system gets started with FIPS mode enabled, Veritas Resiliency Platform operates in the FIPS compliant mode.

You can enable the FIPS mode for a Replication Gateway during the bootstrap process:

See “[Configuring the Replication Gateways](#)” on page 445.

You can enable, disable, or view the status of the FIPS mode by using the `manage > fips` option of klish menu:

Note: You need to keep similar FIPS setting for the paired Replication Gateways.

Password policies for Resiliency Platform virtual appliance

To access Resiliency Platform virtual appliances, you need to set a password for the admin user. Following is the list of rules to set the password:

- Must be at least 8 characters long.
- Must contain at least one uppercase letter (A-Z), one lowercase letter (a-z), one numeric (0-9), and one special character such as @&%.

Note: Though special characters are allowed in a password, you cannot use a space, dollar sign (\$) or double quotes (") in a password. However, this is applicable only in case of Resiliency Platform Data Mover.

- Cannot contain the user name or its characters in reversed order.
- Cannot contain same character used consecutively for more than 2 times.
- Cannot contain 5 or more characters from the previous password.
- Cannot be the same as your previous 6 passwords.
- Can be changed after a minimum of 15 days since the last password change. The password can be changed only through klish menu.
- Expires in 90 days. you get an error message when you are not able to login using admin user credentials.
- 7 days before the password expiry date, a warning is provided to change the password. This warning is not displayed in the Resiliency Manager console. You get to see this warning only when you log in to the virtual appliance console or in the SSH session using the admin user credentials.
- Maximum 10 authentication attempts are allowed within a duration of 15 minutes. After this limit, the user ID and password gets blocked for next 60 minutes.
- In case you forget the admin password, you need to contact Veritas support. Veritas support can reset the admin password only if the virtualization technology or cloud environment supports the serial console.

About NAT support in Veritas Resiliency Platform

Network Address Translation (NAT) is a process in which one or more computers inside a private network are assigned a public address. NAT reduces the need for IPv4 public addresses and hides private network address ranges.

Resiliency Platform 3.3 provides support for NAT to enable communication from a private network to an external network. If there is a non-routable network between the source data center and the target data center, then You need to configure NAT only using IPv4 address for communication between Resiliency Platform appliances.

Considerations for configuring NAT

- Resiliency Platform supports NAT only for the communication between Resiliency Platform components deployed in different data centers. These components need to communicate with each other over public IP address. NAT can exist in both the data centers or in any one of the data centers.
- Resiliency Platform does not support NAT for communication between components deployed in the same data center. The components based in the same data center can communicate with each other over private IP address.
- If a communication channel has been setup in a way that private IP address of the Resiliency Platform components is accessible in the other data center, then you need not configure NAT during the bootstrap process. For example, if you have setup VPN for communication between on-premises data center to AWS datacenter, then you need not configure NAT for Resiliency Platform components .
- If there are multiple Resiliency Managers in one data center, then either NAT configuration should be done for all of them or for none of them. A mix of NAT configured Resiliency Managers and non-NAT configured Resiliency Managers within a single data center is not supported.

Following are the scenarios in which NAT configuration is required for Resiliency Platform virtual appliances:

Table 1-37 Scenarios in which NAT configuration is required

Communication between	Description	NAT configuration
Replication Gateway with Peer Replication Gateway	A Replication Gateway can communicate with the peer Gateway only over Public IP.	Required
IMS with Resiliency Manager	If an IMS reports to a Resiliency Manager within the same data center, then it can communicate over private IP.	Not required
	If an IMS reports to a Resiliency Manager deployed in another data center, then it can communicate only over public IP.	Required

Table 1-37 Scenarios in which NAT configuration is required (*continued*)

Communication between	Description	NAT configuration
Resiliency Manager with another Resiliency Manager	If there are multiple Resiliency Managers in one data center and no Resiliency Manager in another data center, then all of the Resiliency Managers within the data center can communicate with each other over private IP.	Not required
	If there are multiple Resiliency Managers in one data center and at least one Resiliency Manager in another data center, then all of the Resiliency Managers need to communicate with each other only over public IP. In this case, you need to enable NAT reflection (Loopback NAT) for the data center where multiple Resiliency Managers are located.	Required

Set up the resiliency domain

Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console. Refer to the below topics:

- See [“Getting started with a new configuration”](#) on page 453.
- See [“Adding an IMS ”](#) on page 456.
- See [“Adding a Replication Gateway”](#) on page 111.
- See [“Adding Google Cloud Platform data center”](#) on page 458.
- See [“Managing user authentication and permissions”](#) on page 463.
- See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Getting started with a new 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.

Table 1-38 Getting Started wizard

Wizard setup	Details
<p>1. Create or Join a Resiliency Domain</p>	<p>For a new Resiliency Platform deployment, select the option Create Resiliency Domain and supply a name for the domain. You can choose whether to allow collection of product usage information.</p> <p>Select Join Resiliency Domain option if you already have a resiliency domain set up and want to add another Resiliency Manager to the existing domain.</p> <p>Refer to See “Adding a Resiliency Manager to an existing resiliency domain” on page 455. Click Continue.</p>
<p>2. Set up the Resiliency Manager</p>	<p>Select Cloud Data Center? if this is a public cloud data center.</p> <p>Enter the geographical location of the data center in Data Center Location.</p> <p>Provide a user friendly name to the data center in Data Center Name.</p> <p>Provide a user friendly name to the Resiliency Manager in Resiliency Manager Name.</p> <p>Default entries are shown if the Resiliency Manager has external Internet access to determine the geographical location.</p> <p>Click Create.</p>
<p>3. Set up Authentication Domain</p>	<p>Optional.</p> <p>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.</p> <p>If you want to assign a different user as Super admin you must first set up an LDAP or Active Directory authentication domain.</p> <p>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.</p>
<p>4. Set up Users and Personas</p>	<p>Optional.</p> <p>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.</p> <p>The user with the Super Admin persona can add other users and groups and assign them personas later using the Settings page.</p>

Table 1-38 Getting Started wizard (*continued*)

Wizard setup	Details
5. Set up Cloud Configuration	Optional. This step is enabled only if you select Cloud Data Center in the step 2. You can skip this step and add the cloud configuration later from the console. The wizard verifies the information you enter and notifies you if the information is invalid.
6. Finish Getting Started	You exit the Getting Started wizard. The Dashboard is displayed and from the Settings page you can complete any steps that you have skipped.

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 as well as 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 and the Admin login credentials for an existing Resiliency Manager virtual appliance in the resiliency domain.
 - In case of multiple Resiliency Managers in a data center, ensure the following:
 - All the existing Resiliency Managers must be online in the data center where you plan to add a new Resiliency Manager.
 - For a cloud data center, it is recommended to have a minimum of three Resiliency Managers if you want to have multiple Resiliency Managers in the data center.
- 2 Log in to the web console on the new Resiliency Manager. The Getting Started wizard is displayed.

- 3** In **Create or Join a Resiliency Domain**, select **Join resiliency domain**.
Enter the fully qualified host name of a Resiliency Manager in the domain you want to join, user name, and password for the Resiliency Manager, and click **Verify**.
- 4** In **Set up Resiliency Manager**, specify the data center location, the data center friendly name, and Resiliency Manager friendly name.
- 5** Click **Confirm & Continue**.
- 6** After the host name, user name, and password has been verified, the Resiliency Domain Name appears automatically. Select one of the following data centers and click **Continue**.
 - **Create new data center**
 - **Select from existing data center**
- 7** 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 456.
- 8** 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.

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

- If the IMS is in AWS or Google Cloud Platform (GCP) cloud data center, ensure that the IMS has an IAM role with all the required permissions attached to it.
AWS: See “Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway” on page 389.
GCP: See “API permissions required for Google Cloud Platform” on page 460.
 - Information needed for adding the IMS:
Provide hostname.
The admin user credentials for the IMS virtual appliance. This information is optional and you need to enter only if the server is directly accessible. If the server is not directly accessible, you can still initiate the process of adding an IMS by entering only the data center, friendly name, and FQDN/IP address. In this case, you get a registration URL which you have to use after logging in to the virtual appliance console of the IMS that you want to add and then the IMS is added to the data center.
 - Ensure that the IP address and hostname of the IMS gets resolved from the Resiliency Manager.
- 2** Navigate to **Settings** (menu bar) > **Infrastructure** > **Details View** and then Select **+ Infrastructure Management Server**.
- You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.
- 3** In **Add Infrastructure Management Server**, enter the information for the IMS and submit.
- Tips:
- You can select from a list of existing data centers or add a new data center.
- To specify a new data center, select **New** in the **Data Center** field, then specify the location and name. When entering the location, enter a form of location identifier, such as city, and the location list will populate with potential matches for you to select.
 - Enter a friendly name for the IMS.
 - Enter the FQDN or IP address of the server.
 - Enter the user name and the password. These two are optional information that you need to enter only if the IMS is directly accessible. If you provide this information, the IMS is immediately added to the data center. If you do not provide the username and password of the IMS, a registration URL is displayed on the screen. This URL is valid only for 30 minutes. If the URL expires, you need to regenerate the registration URL to complete the process.

Copy the URL string and then log in to the virtual appliance console of the IMS. In the klish menu, run the following command:

```
manage > configure ims_register
```

See [“Klish menu options for IMS”](#) on page 598.

You are prompted to provide the IMS registration URL. Enter the URL that you had obtained after initiating the process from Resiliency Manager console.

4 Verify that the IMS is successfully added.

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

Note: When the IMS is connected to a RM in its own datacenter, that IMS is displayed alongside the RM in the same row on the UI. However, when the IMS is connected to a RM in another datacenter, that IMS is displayed in a row below the RM.

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.

Adding a Replication Gateway

If you plan to use Resiliency Platform Data Mover for replication of data in your environment, you need to deploy and configure at least one Replication Gateway in your source as well as target (recovery) data center. After you have deployed and configured the gateway on each data center, you need to associate the gateway with the IMS on that data center.

For a Replication Gateway to be used for VMware based in cloud, the replication gateway has to be deployed inside the VMware Software Defined Data Center (SDDC) in order for it to have access to the ESX cluster and datastores to be used in recovery.

To add a Replication Gateway

1 Navigate



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

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

2 Click **Data Mover** card. Under **Data Mover** , click **+ Add Replication Gateway**.

3 In the **Replication Gateway** panel, enter the following details, and click **Submit**.

- Select the Infrastructure Management Server (IMS) to which you want to add the Replication Gateway.
- IP address or DNS name of the server.
- User name and password.

You can use the following icons for entering details of multiple Replication Gateways simultaneously or for deleting a particular row from the table:

■ **Icon Task**



To add a blank table row.



To copy the details of the selected table row. You can edit the details of the newly added row.



To import the information from a CSV file. Click **Browse** to select the text file and then click **Load host details**.



To delete a row.

Adding Google Cloud Platform data center

To access the cloud resources, you need to configure the cloud configuration. This information is used to validate the cloud configuration. When you add an Google Cloud Platform cloud data center, you enter a few details for the data center such as Project ID and Service Account like Client Emails, Private key if not autodetected. You need to enter details for region and bucket name also.

To add Google Cloud Platform data center

- 1 Prerequisite:
 - Ensure that the Service Account which is used to add the Google Cloud Platform data center has the required API permissions.
 - Ensure that the subnet of the NIC with which Resiliency Manager can communicate with GCP having private Google Access should be turned on or should have internet access.
 - Ensure to enable Compute Engine, Cloud KMS, and Cloud Storage APIs.See [“API permissions required for Google Cloud Platform”](#) on page 460.
- 2 Navigate to **Settings (menu bar) > Infrastructure > Details view**.
- 3 Select **Datacenter +**.
- 4 Enter the geographical location and the name of the data center.
- 5 Select **Is Cloud Datacenter** and then select **Google Cloud Platform** as cloud type. Click **Next**.
- 6 In the **Google Cloud Platform configurations** window, on **Service Account details for data center** panel select any one option to enter the following information:

- **Continue with associated account**
- **Provide another service account details**

If you select **Continue with associated account** option, provide following information:

- Enter the configuration name.
Provide a user friendly name to the configuration.
- **Project ID:** This is (auto-detected)
- **Service Account:**This is (auto-detected)
- **Region:** Select from the dropdown.
- **Bucket:** Select from the dropdown.
- You can select or unselect the **Discover the Shared VPC Resources** checkbox.

Note: If this checkbox is checked, at least one subnet from same region needs to be shared from host project with configured project. You need to configure few permissions for this checkbox. Refer [API permissions required for Google Cloud Platform](#)

If you select **Provide another service account details** option, provide following information:

- Enter the configuration name.
Provide a user friendly name to the configuration.
- **Project ID:** Enter the Project ID.
- **Client Email:** Enter the Client Email.
- **Region:** Select from the dropdown.
- **Bucket:**Select from the dropdown.
- You can select or unselect the **Discover the Shared VPC Resources** checkbox.

Note: If this checkbox is checked, at least one subnet from same region needs to be shared from host project with configured project. You need to configure few permissions for this checkbox. Refer [API permissions required for Google Cloud Platform](#)

- Enter Private Key. Click **Verify** to validate the private key.
This key can be extracted from Service account json. Refer [Creating and managing service account keys](#)

7 Click **Submit** to complete the configuration.

[Managing cloud configurations](#)

[Refreshing cloud data center](#)

API permissions required for Google Cloud Platform

Following are the permissions required for the roles that you need to create for Resiliency Manager and IMS for recovery to Google Cloud Platform data center. If you are deploying your appliances through Marketplace, then refer to

[API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace.](#)

Table 1-39 API Permissions required for role for Resiliency Manager

Service name	Permissions
compute	compute.regions.list
storage	storage.buckets.list

Table 1-40 API Permissions required for role for IMS

Service name	Permission
cloudkms	cloudkms.cryptoKeyVersions.list
	cloudkms.cryptoKeys.list
	cloudkms.keyRings.list

Table 1-40 API Permissions required for role for IMS (*continued*)

Service name	Permission
compute	compute.addresses.list
	compute.diskTypes.list
	compute.disks.create
	compute.disks.createSnapshot
	compute.disks.delete
	compute.disks.list
	compute.disks.use
	compute.disks.get
	compute.firewalls.list
	compute.globalOperations.list
	compute.images.create
	compute.images.get
	compute.images.useReadOnly
	compute.instances.attachDisk
	compute.instances.create
	compute.instances.delete
	compute.instances.detachDisk
	compute.instances.get
	compute.instances.list
	compute.instances.setMetadata
	compute.instances.setTags
	compute.instances.start
	compute.instances.stop
compute.machineTypes.list	
compute.networks.list	

Table 1-40 API Permissions required for role for IMS (*continued*)

Service name	Permission
	compute.projects.get
	compute.regionOperations.list
	compute.snapshots.create
	compute.snapshots.delete
	compute.snapshots.list
	compute.snapshots.useReadOnly
	compute.subnetworks.list
	compute.subnetworks.use
	compute.zoneOperations.list
	compute.zones.list
	compute.addresses.useInternal
storage	storage.objects.create
	storage.objects.get

Table 1-41 Permissions required to discover Shared VPC for a role for IMS on the host project.

Service name	Permission
compute	compute.networks.list
	compute.firewalls.list
	compute.addresses.list

Note: To share subnets from the host project with configured project, you need to add the Service Account associated with IMS as **Principal** and provide a **Compute Network User role** on the shared subnets.

Managing cloud configurations

See [“Adding AWS cloud data center”](#) on page 118.

See [“Adding Azure cloud data center”](#) on page 119.

See [“Adding Orange Recovery Engine data center”](#) on page 120.

See [“Adding Google Cloud Platform data center”](#) on page 458.

See [“Adding vCloud Director cloud data center”](#) on page 130.

See [“Refreshing cloud data center”](#) on page 133.

See [“Removing the cloud configuration”](#) on page 134.

Adding AWS cloud data center

To access the Cloud resources, you need to configure the cloud credentials. These credentials are used to interface with Cloud APIs. When you add an AWS cloud data center, you enter a few details for the data center such as region and S3 bucket name. Resiliency Manager validates these information using a role that has certain permissions attached to it.

To add AWS cloud data center

1 Prerequisite

Ensure that the Resiliency Manager has an IAM role with all the required permissions attached to it.

See [“Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway”](#) on page 389.

2 Navigate



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

3 Select **Datacenter +**.

4 Enter the geographical location and the name of the data center.

5 Select **Is Cloud Datacenter** and then select AWS cloud type.

6 In the **AWS Customization** panel, enter the following information:

- Enter the Configuration name.
Provide a user-friendly name to the configuration.
- Select the region.
- Enter the name of the S3 bucket that already exists in the region and is accessible using the IAM role attached to the Resiliency Manager instance.

7 Click **Submit** to complete the configuration.

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

See [“Managing cloud configurations”](#) on page 117.

See [“Editing cloud configuration”](#) on page 132.

See [“Refreshing cloud data center”](#) on page 133.

Adding Azure cloud data center

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

Cloud APIs for the data center in Azure perform the following operations:

- Discover virtual machines using near real-time discovery (NRT) , volumes, storage accounts, resource groups, virtual networks, subnets, disks, network security groups, standards or virtual machine size.
- Invoke operations on the cloud resources such as provision virtual machines, attach or detach disks.

Note:

To add Azure cloud data center

1 Navigate



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

2 Select **Datacenter +**.

3 Enter the geographical location and the name of the data center.

4 Select **Is Cloud Datacenter** and then select Azure cloud type.

5 In the **Azure Customization** panel, enter the following information:

- Enter the configuration name.
Provide a user-friendly name to the configuration.
- Select the endpoint.
- Enter the Subscription ID.

- Enter the Tenant ID / Directory ID.
- Enter the Application ID / Client ID.
- Enter the Client Secret Key. Make sure you to use only Client Secret based credentials for applications.

For information on creating an Application ID and Client Secret key, refer to [Microsoft documentation](#). The application is required to communicate with Azure programmatically. Client Secret Key is sometimes also referred as *Authentication key* or *Password*.

- 6 If you want to use an existing configuration, select the check box which then auto populates all the information.
- 7 Click **Verify** to verify the configuration.
- 8 Select a location.
- 9 Select a storage account and enter the container name.
Storage accounts of the type *sku Premium* are not supported.
- 10 Click **Submit** to complete the configuration.

Note: 1. From the console, you can add multiple IMSs to a single cloud data center, but this configuration is not yet supported in Veritas Resiliency Platform.

Note: 2. Authentication of Active Directory Federation Services (ADFS) based configuration is not supported in Azure cloud configuration.

See [“Managing cloud configurations”](#) on page 117.

See [“Editing cloud configuration”](#) on page 132.

See [“Refreshing cloud data center”](#) on page 133.

Adding Orange Recovery Engine data center

To access Orange Recovery Engine resources, you need to configure the cloud credentials. These credentials are used to connect to Orange Recovery Engine using Orange Recovery Engine APIs. You can configure cloud credentials as part of the initial getting started wizard or add the cloud configuration later from the **Infrastructure Settings** or by editing the cloud configuration. This configuration is then automatically pushed in Infrastructure Management Server (IMS).

Orange Recovery Engine provided cloud APIs are used to perform the following operations on the cloud resources:

- Discover virtual machines, disks, virtual networks, subnets, disks, network security groups, server type or flavors.
- Invoke operations on the cloud resources such as provision virtual machines, attach or detach disks.

To add Orange Recovery Engine data center

1 Prerequisites:

Ensure that the user which is used to add the Orange Recovery Engine data center has the required API permissions.

See [“API permissions required for Orange Recovery Engine”](#) on page 122.

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

3 Select **Datacenter +**.

4 Enter the geographical location and the name of the data center.

5 Select **Is Cloud Datacenter** and then select Orange Recovery Engine as cloud type. Click **Next**.

6 In the **Orange Recovery Engine Configuration** panel, enter the following information:

- Enter the configuration name.
Provide a user friendly name to the configuration.
- Select the region.
- Configuration URL is auto-populated.
- Enter the username.
- Enter the API password.
- Enter the project.
- Enter the account name. Click **Next**.
- Click **Finish** to complete the configuration.

Note: Although you can add multiple IMSs to a single cloud data center using the console, provided that all other IMSs in that data center are disconnected.

See [“Managing cloud configurations”](#) on page 117.

See [“Editing cloud configuration”](#) on page 132.

See [“Refreshing cloud data center”](#) on page 133.

API permissions required for Orange Recovery Engine

To create a user in Orange Recovery Engine, you should have access to domain admin account.

User group is a group of users who share set of permissions. When a user group is assigned to a user, it has all the permissions that are assigned to that user group.

User groups are associated with project with some policies. A policy consists of set of permission and these policies are associated with multiple projects.

There are some policies which are default and some policies which are custom in which you can create your own policy for a project. You cannot modify default policies, but you can change custom policies.

A project in a region is associated with multiple user groups which can be assigned to a user.

For example: A user USER1 is associated with a user group UG1 on project P1 with say tenant-admin policy.

See [“Example of a policy statement for Orange Recovery Engine”](#) on page 124.

Following are the API permissions required for the user to discover the assets in the data center and perform operations on these assets.

Table 1-42 API Permissions required for discovery of assets and perform operations:

Service name	Permission
ecs	ecs:serverInterfaces:get ecs:serverInterfaces:use ecs:serverVolumeAttachments:create ecs:serverVolumeAttachments:delete ecs:serverVolumeAttachments:list ecs:serverVolumes:use ecs:servers.stop ecs:servers:create ecs:servers:delete ecs:servers:get ecs:servers:list ecs:servers:start ecs:availabilityZones:list ecs:flavors:get ecs:securityGroups:use ecs:cloudServers:detachVolume ecs:cloudServers:attachVolume

Table 1-42 API Permissions required for discovery of assets and perform operations: *(continued)*

Service name	Permission
evs	evs:volumes:attach evs:volumes:unmanage evs:volumes:create evs:volumes:delete evs:volumes:detach evs:volumes:get evs:volumes:list evs:volumes:manage evs:volumes:update evs:snapshots:create evs:snapshots:delete evs:snapshots:get evs:types:get
ims	ims:images:create ims:images:get ims:images:list ims:images:update ims:images:upload
vpc	vpc:networks:get vpc:ports:create vpc:ports:delete vpc:ports:get vpc:routers:get vpc:subnets:get vpc:securityGroups:get

Example of a policy statement for Orange Recovery Engine

Following is an example of policy statement that you can use to manage the permissions for the users. This policy assigns the following permissions:

```
{
  "Version": "1.1",
  "Statement": [
    {
      "Action": [
        "ecs:*:*",
        "evs:*:get",
        "evs:*:list",
        "evs:volumes:create",
        "evs:volumes:delete",
        "evs:volumes:attach",
        "evs:volumes:detach",
        "evs:volumes:manage",
        "evs:volumes:update",
        "evs:volumes:uploadImage",
        "evs:snapshots:create",
        "vpc:*:get",
        "vpc:*:list",
        "vpc:ports:get",
        "vpc:ports:create",
        "vpc:securityGroups:get",
        "vpc:floatingIps:get",
        "ims:images:create",
        "ims:images:get",
        "ims:images:list",
        "ims:images:upload"
      ],
      "Effect": "Allow"
    }
  ]
}
```

Adding Google Cloud Platform data center

To access the cloud resources, you need to configure the cloud configuration. This information is used to validate the cloud configuration. When you add an Google Cloud Platform cloud data center, you enter a few details for the data center such as Project ID and Service Account like Client Emails, Private key if not autodetected. You need to enter details for region and bucket name also.

To add Google Cloud Platform data center

1 Prerequisite:

- Ensure that the Service Account which is used to add the Google Cloud Platform data center has the required API permissions.

- Ensure that the subnet of the NIC with which Resiliency Manager can communicate with GCP having private Google Access should be turned on or should have internet access.
- Ensure to enable Compute Engine, Cloud KMS, and Cloud Storage APIs.

See [“API permissions required for Google Cloud Platform”](#) on page 460.

- 2 Navigate to **Settings (menu bar) > Infrastructure > Details view**.
- 3 Select **Datacenter +**.
- 4 Enter the geographical location and the name of the data center.
- 5 Select **Is Cloud Datacenter** and then select **Google Cloud Platform** as cloud type. Click **Next**.
- 6 In the **Google Cloud Platform configurations** window, on **Service Account details for data center** panel select any one option to enter the following information:

- **Continue with associated account**
- **Provide another service account details**

If you select **Continue with associated account** option, provide following information:

- Enter the configuration name.
Provide a user friendly name to the configuration.
- **Project ID:** This is (auto-detected)
- **Service Account:**This is (auto-detected)
- **Region:** Select from the dropdown.
- **Bucket:** Select from the dropdown.
- You can select or unselect the **Discover the Shared VPC Resources** checkbox.

Note: If this checkbox is checked, at least one subnet from same region needs to be shared from host project with configured project. You need to configure few permissions for this checkbox. Refer [API permissions required for Google Cloud Platform](#)

If you select **Provide another service account details** option, provide following information:

- Enter the configuration name.

Provide a user friendly name to the configuration.

- **Project ID:** Enter the Project ID.
- **Client Email:** Enter the Client Email.
- **Region:** Select from the dropdown.
- **Bucket:**Select from the dropdown.
- You can select or unselect the **Discover the Shared VPC Resources** checkbox.

Note: If this checkbox is checked, at least one subnet from same region needs to be shared from host project with configured project. You need to configure few permissions for this checkbox. Refer [API permissions required for Google Cloud Platform](#)

- Enter Private Key. Click **Verify** to validate the private key. This key can be extracted from Service account json. Refer [Creating and managing service account keys](#)

7 Click **Submit** to complete the configuration.

[Managing cloud configurations](#)

[Refreshing cloud data center](#)

API permissions required for Google Cloud Platform

Following are the permissions required for the roles that you need to create for Resiliency Manager and IMS for recovery to Google Cloud Platform data center. If you are deploying your appliances through Marketplace, then refer to

[API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace.](#)

Table 1-43 API Permissions required for role for Resiliency Manager

Service name	Permissions
compute	compute.regions.list
storage	storage.buckets.list

Table 1-44 API Permissions required for role for IMS

Service name	Permission
cloudkms	cloudkms.cryptoKeyVersions.list
	cloudkms.cryptoKeys.list
	cloudkms.keyRings.list

Table 1-44 API Permissions required for role for IMS (*continued*)

Service name	Permission
compute	compute.addresses.list
	compute.diskTypes.list
	compute.disks.create
	compute.disks.createSnapshot
	compute.disks.delete
	compute.disks.list
	compute.disks.use
	compute.disks.get
	compute.firewalls.list
	compute.globalOperations.list
	compute.images.create
	compute.images.get
	compute.images.useReadOnly
	compute.instances.attachDisk
	compute.instances.create
	compute.instances.delete
	compute.instances.detachDisk
	compute.instances.get
	compute.instances.list
	compute.instances.setMetadata
	compute.instances.setTags
	compute.instances.start
	compute.instances.stop
	compute.machineTypes.list
compute.networks.list	

Table 1-44 API Permissions required for role for IMS (*continued*)

Service name	Permission
	compute.projects.get
	compute.regionOperations.list
	compute.snapshots.create
	compute.snapshots.delete
	compute.snapshots.list
	compute.snapshots.useReadOnly
	compute.subnetworks.list
	compute.subnetworks.use
	compute.zoneOperations.list
	compute.zones.list
	compute.addresses.useInternal
storage	storage.objects.create
	storage.objects.get

Table 1-45 Permissions required to discover Shared VPC for a role for IMS on the host project.

Service name	Permission
compute	compute.networks.list
	compute.firewalls.list
	compute.addresses.list

Note: To share subnets from the host project with configured project, you need to add the Service Account associated with IMS as **Principal** and provide a **Compute Network User role** on the shared subnets.

Adding vCloud Director cloud data center

To access the Cloud resources, you need to configure the cloud credentials. These credentials are used to interface with Cloud APIs. You can configure cloud

credentials as part of the initial getting started experience or add the cloud configuration later to the Infrastructure Management Server (IMS) on the cloud.

Cloud APIs for the data center in vCloud Director perform the following operations:

- Discover vApps and virtual machines

If there are multiple virtual data centers associated with the Organization, you need to select the virtual data center which you want to use for the cloud configuration. You can add only one virtual data center at a time. Using the add data center operation you can add the other virtual data centers later. Ensure that you deploy and configure the Infrastructure Management Server (IMS) and the replication gateway in each of the virtual data center.

To add vCloud Director cloud data center

- 1 Navigate



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

- 2 Select **Datacenter +**.
- 3 Enter the geographical location and the name of the data center.
- 4 Select **Is Cloud Datacenter** and then select vCloud cloud type.
- 5 In the **vCloud Organization Configuration: Probe** panel, enter the following information:
 - Enter vCloud URL.
 - Enter Organization name.
 - Enter the username for the Organization. This user must have organization administrator role assigned.
 - Enter password for the user.
 - Enter the port number if you want to change the default port.
 - Click **Next** after filling in the details. If there are multiple virtual data centers associated with the Organization, you need to select the virtual data center which you want to use for cloud configuration.
- 6 Click **Submit** to complete the configuration.

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

For secure communication, refer

See [“Managing cloud configurations”](#) on page 117.

See [“Editing cloud configuration”](#) on page 132.

See [“Refreshing cloud data center”](#) on page 133.

Editing cloud configuration

Using the Veritas Resiliency Platform console, you can edit the configuration of a cloud server.

To edit the configuration of a cloud server

- 1 Navigate



Settings (menu bar)

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

- 2 Click the vertical ellipses next to the cloud data center, and click **Edit Datacenter**.
- 3 In the **Edit Datacenter** wizard panel, do one of the following:

For AWS cloud configuration

After you complete adding AWS cloud configuration for the first time, a new virtual private cloud (VPC) is added to the AWS region, but the VPC and its components are not discovered and available in the console. You need to edit the cloud configuration and then save the configuration to make the newly added VPC discovered.

For AWS cloud configuration, you can only edit the configuration name.

You can not edit the region.

For Azure cloud configuration

You can edit the following:

- Application ID
- Client secret key

Click **Verify** to edit the below information.

- Account name
- Container name

Location cannot be edited if the cloud server configuration is already saved on the Infrastructure Management Server (IMS).

Storage accounts of the type *sku Premium* are not supported.

For vCloud Director cloud configuration

You can edit the following:

- Username
- Password
- Port

4 Click **Submit** to complete the configuration.

See [“Managing cloud configurations”](#) on page 117.

Refreshing cloud data center

Using the Veritas Resiliency Platform console, you can refresh the cloud data center.

If your data center is in AWS, this operation re-discovers all the cloud-based objects such as regions, zones, security groups, flavors, virtual machines in the cloud.

If your data center is in vCloud Director, this operation re-discovers the vApps, virtual machines and the cloud network.

If your data center is in Azure, this operation re-discovers resource groups, virtual networks, subnets, virtual machines, disks, network security groups, standards or virtual machine size.

To refresh cloud data center

- 1 Navigate



Settings (menu bar)

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

- 2 Click the vertical ellipses next to the cloud data center, and click **Refresh Cloud Discovery**.
- 3 In the **Refresh Cloud Discovery** wizard panel, click **Next**.
- 4 Click **Submit** to close the wizard.

See [“Managing cloud configurations”](#) on page 117.

Removing the cloud configuration

Using the Veritas Resiliency Platform console, you can disconnect the cloud data center from the Infrastructure Management Server (IMS). This operation removes all the discovered objects such as regions, zones, security groups, flavors, virtual machines.

To remove a cloud configuration

- 1 Navigate



Settings (menu bar)

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

- 2 Click the vertical ellipses next to the cloud data center, and click **Remove Datacenter**.
- 3 In the **Remove Datacenter** wizard panel, click **Submit**.

See [“Managing cloud configurations”](#) on page 117.

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 1-46 Process for setting up user authentication and permissions

Task	Details
Configure authentication domains	<p>You can add multiple authentication domains.</p> <p>See “About user authentication in the web console” on page 464.</p> <p>See “Configuring authentication domains” on page 469.</p> <p>See “Unconfiguring authentication domains” on page 474.</p>
Configure user groups and users	<p>Once you configure an authentication domain, you can configure user groups or users for Resiliency Platform from that authentication domain.</p> <p>See “Configuring user groups and users” on page 475.</p>
Assign permissions to groups and users	<p>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.</p> <p>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.</p> <p>See “About user permissions in the web console” on page 464.</p> <p>See “Predefined personas” on page 465.</p> <p>See “About limiting object scope for personas” on page 482.</p> <p>See “Assigning permissions to user groups and users” on page 476.</p> <p>You can also create custom personas.</p> <p>See “Adding custom personas” on page 478.</p> <p>See “Predefined jobs that can be used for custom personas” on page 479.</p>
Configure Windows global user	<p>To customize the static IP of Windows guest virtual machines in the VMware environment, you need to provide the administrator user name and password to log on to the Windows virtual machines.</p> <p>See “Configuring Windows global user” on page 482.</p>

About user authentication in the web console

By default, the Admin user of the Veritas Resiliency Platform virtual appliance can log in to the web console with access to all views and operations.

The Admin user can configure authentication domains from external identity providers such as Active Directory (AD) and LDAP.

Once an authentication domain is configured, the Admin user can configure user groups and users for Resiliency Platform from that domain. These users can log in to the console with their domain login credentials.

All users and groups that are configured for Resiliency Platform have permission by default to view everything in the web console but not to perform any operations. Permissions for operations must be assigned separately by assigning the appropriate personas to users and groups.

It is recommended not to remove the default Resiliency Platform users or reduce the permissions of the default Resiliency Platform users.

If you change the password of a user who was configured to logon to the domain, you need to edit the configured domain and enter the new password for the user.

See [“Editing authentication domains”](#) on page 475.

See [“Managing user authentication and permissions”](#) on page 463.

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	<p>A role that has access to a predefined set of jobs (operations).</p> <p>The product comes with a set of predefined personas.</p> <p>See “Predefined personas” on page 465.</p> <p>You can also add custom personas.</p> <p>See “Adding custom personas” on page 478.</p> <p>See “Predefined jobs that can be used for custom personas” on page 479.</p> <p>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.</p>
---------	---

Job	<p>A type of task (operation) that a user can perform.</p> <p>Examples:</p> <ul style="list-style-type: none"> Manage resiliency groups Manage assets Perform disaster recovery of resiliency groups
Object types and scope	<p>Each job can be performed on certain types of Resiliency Platform objects. Types of objects include data centers, resiliency groups, and virtual business services.</p> <p>When you assign a persona to a user or group, you define the scope of some jobs by selecting from available objects. For some jobs, the scope is the resiliency domain, which would be the entire scope of the product deployment.</p>

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 463.

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 1-47 Predefined personas and jobs

Persona	Description and scope	Jobs
Super admin	<p>Can perform all operations on all objects in resiliency domain.</p> <p>Scope: Resiliency domain.</p>	<ul style="list-style-type: none"> ■ All jobs ■ All objects in resiliency domain

Table 1-47 Predefined personas and jobs (*continued*)

Persona	Description and scope	Jobs
Resiliency Platform admin	<p>Manage Resiliency Managers and Infrastructure Management Servers (IMs) and data centers.</p> <p>Manage assets.</p> <p>Manage user security settings and other product settings.</p> <p>Manage product updates.</p> <p>Scope: Resiliency domain.</p>	<ul style="list-style-type: none"> ■ Manage enclosure assets ■ Manage server deployments ■ Manage product updates ■ Manage virtualization assets ■ Manage application host ■ Manage application cluster assets ■ Manage DR settings ■ Manage user security settings ■ Manage copy manager assets ■ Manage product settings ■ Manage access profiles ■ Manage cloud assets ■ Manage data mover assets
Resiliency Platform Deployment admin	<p>Manage Resiliency Managers and Infrastructure Management Servers (IMs).</p> <p>Can add an IMS to an existing data center.</p> <p>Manage product updates.</p> <p>Scope: Resiliency domain.</p>	<ul style="list-style-type: none"> ■ Manage server deployments ■ Manage product updates

Table 1-47 Predefined personas and jobs (*continued*)

Persona	Description and scope	Jobs
Data Center admin	<p>Manage infrastructure pairing and manage assets of specified types.</p> <p>Scope: Specified data center.</p>	<ul style="list-style-type: none"> ■ Execute custom scripts ■ Manage cloud assets ■ Manage enclosure assets ■ Manage copy manager assets ■ Manage application cluster assets ■ Manage data mover assets ■ Manage DR settings ■ Manage application host ■ Manage virtualization assets ■ Manage access profiles
Resiliency Domain admin	<p>Create, update, and delete resiliency groups, virtual business services (VBSs), and resiliency plans and templates.</p> <p>Start/stop all resiliency groups and VBSs.</p> <p>Configure all resiliency groups for disaster recovery (DR).</p> <p>Perform rehearsal and DR operations: migrate, recover.</p> <p>Create, update, and delete resiliency plans and templates.</p> <p>Manage disaster recovery network settings.</p> <p>Scope: Resiliency domain.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Manage resiliency plans ■ Manage virtual business services ■ Manage resiliency plan templates ■ Manage resiliency groups ■ Recover resiliency group ■ Execute custom scripts ■ Rehearse resiliency group

Table 1-47 Predefined personas and jobs (*continued*)

Persona	Description and scope	Jobs
Resiliency Group admin	<p>Update and delete specified resiliency groups.</p> <p>Start/stop specified resiliency groups.</p> <p>Start/stop VBSs as long as the VBS contains only the specified resiliency groups.</p> <p>Scope: Specified resiliency groups.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Manage resiliency groups
Resiliency Group operator	<p>Start/stop specified resiliency groups.</p> <p>Start/stop VBSs as long as the VBS contains only the specified resiliency groups.</p> <p>Scope: Specified resiliency groups.</p>	Start/stop resiliency groups
VBS admin	<p>Create, update, and delete all virtual business services (VBSs).</p> <p>Start/stop all resiliency groups and VBSs.</p> <p>Scope: Resiliency domain.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Manage virtual business services
Resiliency Group Recovery admin	<p>Manage and perform disaster recovery of resiliency groups</p> <p>Start/stop specified resiliency groups.</p> <p>Start/stop or perform DR operations on VBSs as long as the VBS contains only the specified resiliency groups.</p> <p>Scope: Specified resiliency groups.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Manage resiliency groups ■ Recover resiliency groups ■ Rehearse resiliency groups

Table 1-47 Predefined personas and jobs (*continued*)

Persona	Description and scope	Jobs
Resiliency Group Recovery operator	<p>Start/stop specified resiliency groups.</p> <p>Perform disaster recovery on specified resiliency groups.</p> <p>Start/stop or perform DR operations on VBSs as long as the VBS contains only the specified resiliency groups.</p> <p>Scope: Specified resiliency groups.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Recover resiliency groups ■ Rehearse resiliency groups
Guest	<p>View all information in console.</p> <p>Assigned by default when user or group is configured for Resiliency Platform.</p> <p>Scope: Resiliency domain</p>	View all information
Resiliency Platform Assets admin	<p>Manage all assets such as enclosure, application, application cluster assets, virtualization, data mover, and cloud.</p> <p>Scope: Resiliency domain</p>	<ul style="list-style-type: none"> ■ Manage enclosure assets ■ Manage virtualization assets ■ Manage application host ■ Manage application cluster assets ■ Manage copy manager assets ■ Manage access profiles ■ Manage cloud assets ■ Manage data mover assets

See [“Managing user authentication and permissions”](#) on page 463.

Configuring authentication domains

By default, the Admin user on the Veritas Resiliency Platform virtual appliance can log in to the Resiliency Platform web console with access to all views and operations. The Admin user can configure authentication domains for Resiliency Platform from external identity providers so that other users can be authenticated for access to the console.

[To configure authentication domains](#)

[To edit authentication domains](#)

To configure authentication domains

1 Prerequisites

The fully qualified domain name (FQDN) or IP address and credentials for the LDAP/AD servers in the authentication domain. If you are configuring with IPv6 address, specify the hostname and not the IP address.

2 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Domains**

Note: You can also configure an authentication domain from the Getting Started wizard after setting up the Resiliency Manager and resiliency domain.

3 Click **+ Configure Domain**.

4 Select a data center and under **Specify server information for each data center**, enter the information for the server at that data center.

Repeat this step for other data centers in the authentication domain. When you select a different data center, the server information fields are cleared so that you can enter information for a different server, but the entries for the previous data center are remembered.

Note: If the same server is used for more than one data center, enter the same server information for each data center.

The remaining fields on the page apply to all data centers; fill these in as required.

See [“Options for authentication domain configuration”](#) on page 471.

Once you have entered information for all data centers, click **Next**.

5 Verify and complete the configuration:

In the **Domain Name** field, enter a friendly name for the authentication domain. If you configure the login screen to list domains, this name is listed.

Verify that the applicable data centers are listed. To make any changes, click **Back** to return to the previous screen. Once all is complete, click **Submit**.

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

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

To edit authentication domains**1** Navigate to the domain list as described in the procedure to configure authentication domains.**2** Select the authentication domain you want to edit and select the Edit option.

Note the following guidelines when editing:

- To add server information for a new data center, select the applicable data center and fill in the server information.
- To edit existing server information, select the applicable data center.
- To edit other information, you do not need to select each data center; the same information applies to all.
- If a data center no longer uses a separate server, replace the server information for that data center with the information for the server that is being used.
- To remove a data center from the authentication domain, use the Unconfigure option instead of the Edit option.

See [“Unconfiguring authentication domains”](#) on page 474.

See [“Managing user authentication and permissions”](#) on page 463.

Options for authentication domain configuration

The first page of the authentication domain configuration wizard is divided into two areas.

See [Table 1-48](#) on page 472.

See [Table 1-49](#) on page 472.

Server information by data center

You must specify the server information separately for each data center. When you select a different data center the server information fields clear so you can enter

the new information. If the same server is used for multiple data centers, enter the same information for both data centers.

Table 1-48 Server information by data center

Option	Description
Server (Mandatory)	Enter the fully-qualified host name or IP address of the LDAP server. If a secure session is configured with the LDAP server using SSL certificates, you must enter the fully-qualified host name that matches with the fully-qualified host name in the LDAP server certificate.
Port (Mandatory)	Displays the number of the port on which the LDAP server is configured to run. By default, this field displays the port number as 389. You can edit this port number, if required.
Connect using SSL/TLS	Select this check box to use the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) certificates to establish a secure channel between the authentication broker and the LDAP server.
Certificate	Browse to the location of the trusted root CA certificate of the vendor that issued the LDAP server certificate.

Configuration options applicable to all data centers

The remaining fields apply to all data centers; fill these in as required.

Table 1-49 Configuration options applicable to all data centers

Option	Description
The authentication servers require me to log on.	Select this check box if the anonymous operations are disabled on the LDAP server and a bind user ID is required to proceed with configuring the LDAP-based authentication

Table 1-49 Configuration options applicable to all data centers (*continued*)

Option	Description
Bind User Name/DN	<p>Enter the complete Distinguished Name (DN) of the user that is used to bind to the LDAP server.</p> <p>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</p> <p>For example, you can provide the DN as Administrator@enterprise.domainname.com ENTERPRISE\Administrator</p> <p>For RFC 2307 compliant LDAP servers, specify complete bind DN.</p> <p>For example, cn=Manager,dc=vss,dc=veritas,dc=com</p> <p>The LDAP or the AD administrator can provide you the bind user name that you can use.</p>
Password	<p>Enter the password that is assigned to the bind user name that you use.</p>
Query Information:	
User (Mandatory)	<p>Under Query Information, enter the user name based on which the system detects the LDAP server-related settings. Ensure that the user name does not contain any special characters.</p> <p>The system determines the search base based on the user name that you specify in this field.</p>
Group	<p>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.</p> <p>The system determines the search base based on the group name along with the user name that you have specified.</p>

Once you have entered information for all data centers, click **Next**.

The **LDAP standard** panel is introduced to select and discover the LDAP schema for the configuration.

Table 1-50 Configuring attributes for LDAP standards

Option	LDAP standard description	Steps
RFC2307	Resiliency Manager uses LDAP schema RFC2307 standard to populate the required attributes and discover the information from LDAP server.	If you select this option and click Next , on Verify and Configure page the LDAP server uses RFC2307 schema for configuration.
Microsoft Active Directory	Microsoft Active Directory is an Active Directory server which uses LDAP protocol.	If you select this option and click Next , on Verify and Configure page the LDAP server uses Microsoft Active Directory schema for configuration.
Custom	The LDAP server uses the default schema as part of LDAP configuration.	<p>If you select this option, you need to provide following attributes for LDAP server and then click Next. The LDAP server will use the default schema present as part of the LDAP configuration.</p> <ul style="list-style-type: none"> ■ User Name: ■ User ID: ■ User description: ■ Group Name: ■ Group ID: ■ Group description: <p>Click Submit and Done.</p> <p>On Verify and Configure panel, the inputs provided are displayed.</p>

See [“Configuring authentication domains”](#) on page 469.

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 463.

Editing authentication domains

Using Resiliency Platform console, you can edit the configuration of an authentication domain. The newly introduced **Custom attributes panel**, allows to edit the the attributes for the LDAP schema.

To edit an authentication domain

1 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Domains**

2 Right-click the domain and select **Edit**.

3 Edit the values that you want to update and click **Next**.

4 Verify the domain configuration details and click **Submit**.

See [“Managing user authentication and permissions”](#) on page 463.

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 Prerequisite

The names of the user groups or users that you want to configure, as configured in the authentication domain.

2 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Users & Groups**

Note: To edit or remove an existing user or group, right-click the name in the list and select the appropriate option.

3 Click **Configure User or Group**.

4 Select the authentication domain.

5 Type the name of the user group or user. Click **Verify** so that the wizard can verify the name in the domain.

6 You can allow this user to access APIs by selecting the **Allow user to access Resiliency Platform APIs** option.

If you have access to APIs, you can generate an API access key and start using Resiliency Platform APIs. From version 3.5, the user which has persona with job **Manage API access key** can generate or revoke the API access key to other users as well.

Note: This option to provide API access is not applicable for user groups.

7 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 476.

See [“Managing user authentication and permissions”](#) on page 463.

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 464.

To assign permissions to user groups and users

1 Prerequisites

The users and groups must be added to Resiliency Platform before you can assign personas.

See [“Configuring user groups and users”](#) on page 475.

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 463.

Adding custom personas

Veritas Resiliency Platform provides a set of predefined personas with access to predefined jobs.

You can add custom personas by selecting from the predefined jobs.

For example, the predefined persona Resiliency Platform Admin includes the jobs for managing assets, managing security settings, and managing product settings. You could create an "Asset Manager" persona that includes only the managing assets job.

You cannot customize the Super admin persona, which has access to all jobs and all objects in the resiliency domain. You also cannot customize the Guest persona, which can view all information in the console.

To add custom personas

1 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Persona & Jobs > New Persona**

2 In the **New Persona** page, complete the following steps and submit:

- Assign a name and description to the custom persona.
- Select one or more jobs that you want to assign to the persona. The jobs are shown in categories depending on whether the scope is the entire resiliency domain or whether the scope can be customized to specific data centers or assets. Select the job from the appropriate category.

For example, if you want to assign a permission related to managing any resiliency group in the resiliency domain, select **Manage Resiliency Group** under the category of **For entire Resiliency Domain**. But if you want to limit permissions to specific resiliency groups, select **Manage Resiliency Group** under the category **For specific resiliency groups**.

See [“Predefined jobs that can be used for custom personas”](#) on page 479.

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 463.

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 1-51 Jobs for custom personas

Jobs	Description	Scope
View all information	View all information in console.	Resiliency domain
Manage user security settings	Manage authentication domains, users and user groups, personas.	Resiliency domain
Manage product settings	Manage general product settings such as alerts and notifications.	Resiliency domain
Manage server deployments	Edit Resiliency Manager information. Join a Resiliency Manager to a domain or leave a domain. Manage IMSs, including add, remove, edit, reconnect operations.	Resiliency domain
Manage product updates	Perform the operations available from the Product Updates page of the console.	Resiliency domain

Table 1-51 Jobs for custom personas (*continued*)

Jobs	Description	Scope
Manage service objectives	Activate service objectives from templates; manage activated service objectives.	Resiliency domain
Manage assets, by type: <ul style="list-style-type: none"> ■ Manage host assets ■ Manage virtualization assets ■ Manage data mover assets ■ Manage application cluster assets ■ Manage cloud assets ■ Manage copy manager assets ■ Manage enclosure assets ■ Manage access profiles 	Add, edit, or remove specific types of asset infrastructure	Resiliency domain or specific data centers
Manage resiliency groups	Create, update, and delete resiliency groups.	Resiliency domain or specific resiliency groups
Start/stop resiliency groups	Start and stop resiliency groups.	Resiliency domain or specific resiliency groups
Manage virtual business services	Create, update, and delete virtual business services (VBSs).	Resiliency domain or specific VBSs
Manage resiliency plans	Create, update, and delete resiliency plans. Note: The permission to execute a resiliency plan depends on a cumulative check on permissions for individual resiliency groups and VBSs in the plan. See "About limiting object scope for personas" on page 482.	Resiliency domain
Manage resiliency plan templates	Create, update, and delete resiliency plan templates.	Resiliency domain

Table 1-51 Jobs for custom personas (*continued*)

Jobs	Description	Scope
Execute custom scripts	Execute custom scripts as part of resiliency plans.	Resiliency domain or specific data centers
Rehearse resiliency groups	Perform rehearsal and rehearsal cleanup. Note: There is no separate job to perform rehearsal of VBSs. If the assigned scope of this job includes all the resiliency groups in a VBS, Rehearsal operations can be performed on that VBS. See “About limiting object scope for personas” on page 482.	Resiliency domain or specific resiliency groups
Recover resiliency groups	Perform Recovery operations such as migrate, recover, resync. Note: There is no separate job to perform disaster recovery of VBSs. If the assigned scope of this job includes all the resiliency groups in a VBS, DR operations can be performed on that VBS. See “About limiting object scope for personas” on page 482.	Resiliency domain or specific resiliency groups
Manage DR settings	Configure disaster recovery network settings, for example, mapping network settings for disaster recovery or replication gateway pairing.	Resiliency domain or specific data centers

 See [“Predefined personas”](#) on page 465.

 See [“Adding custom personas”](#) on page 478.

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 465.

See [“About resiliency groups with assets”](#) on page 528.

For example, you can assign one user the permission to perform operations on resiliency group RG1 and assign another user the permission to perform operations on RG2.

When planning persona assignments in which you select objects to limit the scope, take the following into account:

- Before you can select the objects such as resiliency groups to limit the scope of operations for a persona, the objects must first be created in Resiliency Platform.
- You need to plan for future maintenance on such limited scope personas. If more objects of that type are added later, you may need to edit existing personas for users or user groups in order to add permissions for the new objects.
- Keep in mind that operations on virtual business services (VBSs) that include multiple resiliency groups will fail unless the user performing the operation has permission for operations on all the resiliency groups in the VBS. The same limitation applies for workflow or resiliency plan operations that include multiple resiliency groups.
For example: a VBS is composed of RG1 and RG2. The operator has permission to perform operations on RG1 but not RG2. If they try to perform operations on the VBS, the operation will fail.

Configuring Windows global user

To perform IP customization on Windows virtual machine in VMware environment, Resiliency Platform requires any one of the following users:

- Domain administrator
- Local user who is part of administrator group on the Windows host where you want to perform IP customization
- Domain user who is part of administrator group on the Windows host where you want to perform IP customization
- UAC settings
- User Account Control: Run all administrators in Admin Approval Mode.

For more information on how to manage the settings, refer to Microsoft documentation.

For Windows Active Directory user, the Active Directory should be common for both, the primary and the recovery data center. The Active Directory should be configured before configuring the Windows Global user.

If a Windows virtual machine is part of a Windows Active Directory, ensure that you log on to the virtual machine at-least once using the Active Directory credentials. This is applicable only for VMware environment and if the recovery is on on-premises data center.

To configure Windows global user

1 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Windows Global User**

2 Click + **Configure User** to configure the user.

3 Select between Active Directory and Workgroup.

4 Enter the administrator user name and password. Click **Verify**.

For Workgroup user, enter user name as workgroupname\username. If the workgroup name is not customized then you can enter only the user name.

5 On successful verification, click **Next** and then **Finish** to submit the information.

See [“Network customization options”](#) on page 523.

Managing settings for alerts and notifications and miscellaneous product settings

See the following topics for information on configuring email and SNMP settings for notifications and reports, setting up rules for event notifications, configuring purge intervals, and changing telemetry settings.

See [“Adding, modifying, or deleting email settings”](#) on page 484.

See [“Adding, modifying, or deleting SNMP settings”](#) on page 486.

See [“Throttling the notifications”](#) on page 486.

See [“Downloading the MIB file”](#) on page 490.

See [“Setting up rules for event notifications”](#) on page 490.

See [“Adding, modifying, or removing Syslog server”](#) on page 491.

See [“Modifying the purge setting for logs and SNMP traps”](#) on page 493.

See [“Enabling or disabling telemetry collection ”](#) on page 493.

See [“Showing domains on login screen ”](#) on page 494.

See [“Downloading log files”](#) on page 494.

Adding, modifying, or deleting email settings

You can configure email settings to be used for different features, such as sending reports or receiving automatic email notifications of events. Veritas Resiliency Platform manages email notifications via Resiliency Managers. When Resiliency Managers are located in different geographical locations, the required email settings are likely different for each location. In that case, you add a separate email configuration for each location. You can send a test email to verify the settings. You can also modify or delete existing email configurations.

To add, modify, or delete email settings

1 Navigate



Settings (menu bar)

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

To add a new email configuration, select **Add Email Configuration**.

To modify or delete an existing one, right-click it and select **Modify** or **Delete**.

2 To add or modify an email configuration, go through the wizard pages and specify the options.

In **Server Information**, specify the following:

Name	Assign a unique name for the email configuration.
Email Server	Valid formats include: Fully Qualified Domain Name (FQDN), IP address, or, if the network handles DNS resolution for host names, a shortened host name. Examples: Host123, Host123.example.com, xxx.yyy.zzz.aaa.
SMTP Port	Enter the SMTP mail server port number. The default is 25.
From Email Address	Enter the email address to be shown as the sender of all the emails that are sent.
Friendly Email Name	Optionally, enter a name to be shown for the From address.
Send To	Enter the email address to which you want to send the email.

3 In **Security**, if you want to implement secure SMTP, select the checkbox and enter the user name and password.

4 In **Select Resiliency Managers**, select a Resiliency Manager in the data center location where these email settings apply.

5 In **Test Email Settings**, enter a valid email address, and enter a subject and message for the test email. Select **Send Test Email** to test your settings.

6 Review the information in the summary and submit

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Adding, modifying, or deleting SNMP settings

When an event takes place, you can configure SNMP traps to be sent. The traps are generated using SNMPv2 version. The community string is set to *public* for the generated traps. Resiliency Platform 10.0 enables support for IPv6 network. You can configure the SNMP traps using IPv6 address. If you want to receive the SNMP traps from Resiliency Platform, you can configure using the below mentioned steps:

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 (IPv4 or IPv6) or name of the host where the SNMP trap console is located. Example: Host123.example.com
SNMP Port	Enter the SNMP port number. The default port for the trap is 162.

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Throttling the notifications

Notifications in Resiliency Platform are raised when changes occur due to new operations that are performed or any configuration settings are changed. Some events in Resiliency Platform are short-lived and are auto-cleared when the event condition is successful. In Resiliency Platform, multiple notifications are raised due to which it is difficult to track the real issues. From version 4.0 of Resiliency Platform, throttling notifications feature is introduced that helps in suppressing the notification

(with specific duration) for some time. The suppressed notification is not seen in the notification list under **Logs > Notification** tab.

For example: If a vCenter server is disconnected, a notification is raised. You can throttle this notification for specific duration and is not seen under the notification list.

If a notification is throttled for some duration, and it gets cleared before the specified duration then this notification is not seen under the notification list under **Logs > Notifications** tab.

For example: When a resiliency group is created, "Replication state synchronizing" notification is raised. If this notification is throttled for 5 minutes, and meanwhile the resiliency group is created and is in ONLINE state before the specified duration, the notification gets cleared and is not seen under the notification list.

Following are the points to remember:

- If you have throttled any notification for a specific duration, it is not seen in the notification list until the specified duration.
- If the condition of the notification is still valid, it shall be displayed in the notification list.
- The associated throttled notifications may get removed if the source is deleted from the Resiliency Platform.
- Notifications can be throttled for infinite duration.

For some events where clearing event is not assigned to a particular source, such events has "indefinite time period" duration set by default.

For example: Risk Service, Reporting Service, Scheduling Service, rg.migrate.success, risk.notify.suppress.risk, vmware.vc.ims.refresh.success, etc.

Using Resiliency Platform console, you can add, edit, and remove the throttling notifications of the respective sources. You can group by object and notification topic from the list. The sources for which the notifications are generated in Resiliency Platform are:

Table 1-52 List of objects for which notifications are generated

List of objects
Discovery Host
Data center
ESX Server
Resiliency Group

Table 1-52 List of objects for which notifications are generated (*continued*)

List of objects
NetBackup Primary
Virtual Business Service
vCenter
Infrastructure Management Server
Cluster
Resiliency Manager
Replication Gateway

Table 1-53 List of services for which notifications are generated

List of services
Scheduling Service
Reporting Service
Risk Service

Disable throttling notifications

If you want to disable the throttling notifications, contact Veritas Support.

Notification Throttling Report

Notification Throttling Report displays all notifications that are currently throttled and are waiting to be raised. To view reports, **Reports menu > Inventory >**

Notification Throttling Report

See [“Viewing reports”](#) on page 580.

More Information:

See [“Add throttle notification”](#) on page 488.

See [“Edit throttle notification”](#) on page 489.

See [“Remove throttle notification”](#) on page 490.

Add throttle notification

To add a throttling notification in Resiliency Platform, perform following steps:

To add a new throttling notification

- 1 Navigate to **Settings (menu bar) > Product Settings > Alerts and Notifications**.
- 2 Click on the **Throttle Notifications** tab.
- 3 To add new throttle notification, click **+ Notifications**.
- 4 In the **Add Notification Topic Throttling Settings** panel, select the following:
 - Select the **Source** from the drop-down and click **Next**.
 - Select the notification source for the respective Source objects and click **Next**.
 - Select the check box for the notification topics for throttling of the Source from the list and click **Next**.
 - Set the duration for throttling the notification from the list and click **Next**. Using **Apply All** option, you can set the same duration to all the notifications at a time.
 - Click **Submit** and **Finish**.

The notification is added under the **Throttle Notifications** tab > Notifications list.

More Information:

See [“Throttling the notifications”](#) on page 486.

Edit throttle notification

To edit the throttling notification in Resiliency Platform, perform following steps:

To edit a new throttling notification

- 1 Navigate to **Settings (menu bar) > Product Settings > Alerts and Notifications**.
- 2 Click on the **Throttle Notifications** tab.
- 3 Select the specific **Source** from the list you wish to edit the notification.
- 4 In the throttling notifications list, select the notification you want to edit by clicking on the vertical ellipse and select **Edit** option. While editing the throttling notification, you can only change the duration of the notification.
- 5 Click **Next** and then **Finish** to save the changes.

See [“Throttling the notifications”](#) on page 486.

Remove throttle notification

To remove the throttling notification from the Resiliency Platform, perform the following steps:

To remove the throttling notification

- 1 Navigate to **Settings (menu bar) > Product Settings > Alerts and Notifications**.
- 2 Click on the **Throttle Notifications**.
- 3 Select the **Source** from the list you wish to remove the notification.
- 4 In the throttling notification list, select the notification you want to remove. Click on the vertical ellipses select **Remove** option. On the **Remove throttle notification** panel, you can **View Details** of the notification before removing it.
- 5 Click **Submit** to remove the throttle notification.

Note: If you are removing the throttling notification then the associated notification are also removed.

See [“Throttling the notifications”](#) on page 486.

Downloading the MIB file

You can download the management information base (MIB) file from the Resiliency Manager console. This MIB file defines the format of Veritas Resiliency Platform SNMP traps.

To download the MIB file

- 1 Navigate



Settings (menu bar)

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

- 2 Select **Download MIB file**.

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.

When a resiliency group is placed in maintenance mode, a notification is generated every 24 hrs. You can configure an SNMP or email alert for these notifications.

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 484.

See [“Adding, modifying, or deleting SNMP settings”](#) on page 486.

2 Navigate



Settings (menu bar)

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

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

To modify or delete an existing rule: Select the **Rules** tab, right-click the rule, and select **Modify** or **Delete**.

3 In **Configure Rule**, enter or modify the following:

Name	Enter a unique name for this rule.
Send emails to	Enter one or more email addresses separated by a comma
Send SNMP traps to	Optional
Select Notifications	Select one or more events that you want to be notified about

4 Select **Submit**.

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

Adding, modifying, or removing Syslog server

You can configure Veritas Resiliency Platform to share the Resiliency Manager logs with Syslog server using the Resiliency Platform console. You can configure Syslog server using the IPv6 address.

You can add multiple Syslog servers for a data center. But if you want to collect the audit logs on selected servers, then you need to add those to the Resiliency Platform first.

To add, modify, or remove Syslog server

1 Navigate



Settings (menu bar)

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

To add a new Syslog server, select **Add Syslog Server**.

To modify or delete an existing server, right-click the required server, and select **Edit** or **Remove**.

2 To add or modify Syslog configuration, specify the following:

Data Center Name	Select the data center whose logs you want to send to the Syslog server. Disabled for modify operation.
Syslog Server IP / name	Enter the Syslog server IP address (IPv4 or IPv6) or the name. Disabled for modify operation.
Port	Enter the port number.
Log Level	Select the log level from the following. <ul style="list-style-type: none"> ■ Critical: To share only the critical logs. ■ Error: To share error and critical logs. ■ Warning: To share warning, error, and critical logs. ■ Informational: To share all the logs.
Send Audit Logs	Select if you want to share audit logs with the Syslog server.
Protocol	UDP is the default protocol.

3 Click **Next** and **Finish** to save the changes.

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Modifying the purge setting for logs and SNMP traps

By default, logs and SNMP traps are retained for two years. You can modify this purge setting.

To modify the purge setting for logs and SNMP traps

1 Navigate



Settings (menu bar)

Under **Product Settings**, click **Miscellaneous**

2 Under **Log Settings**, enter the new value for the purge setting, in months, and save the setting.

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

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.

The telemetry information is uploaded to Veritas telemetry collection servers if the Resiliency Manager has Internet connectivity.

To enable or disable telemetry collection

1 Navigate



Settings (menu bar)

Under **Product Settings**, select **Miscellaneous**

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

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Showing domains on login screen

You can set up the login screen to list the available authentication domains. By default, the domains list is not shown and the user must enter a fully qualified username, for example, `username@domain` or `domain\username`.

To show domains on login

1 Navigate



Settings (menu bar)

Under **Product Settings**, select **Miscellaneous**

2 Under **Login Settings**, select **Show domains list** and save the setting.

See [“Managing user authentication and permissions”](#) on page 463.

Downloading log files

Using the Resiliency Platform console, you can download the Resiliency Managers logs for troubleshooting. These support logs are collected by running the `support > loggather` command using the klish menu.

You can view the list of collected support log files for the selected Resiliency Manager. If you have multiple Resiliency Managers, you need to log on to each Resiliency Manager to view its logs. To download the log files, you must have *Resiliency Platform admin* persona with *Manage product settings* job assigned.

Click **Refresh** to view the latest logs.

To download the log files

1 Navigate



Settings (menu bar)

Under **Product Settings**, select **Miscellaneous**

2 Under **Support logs**, select the log file in the list and download.

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Add asset infrastructure

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

- See [“Adding VMware virtualization servers”](#) on page 167.
- See [“Preparing host for replication”](#) on page 496.

Adding VMware virtualization servers

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

- Information on the vCenter Server
- Information on the ESX servers that the vCenter Server manages
When adding a vCenter Server, you have the option to automatically discover all ESX servers registered to the vCenter Server or select which of the available ESX servers to discover.
- Information on the virtual machines that are configured on the ESX servers

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

You can add a VMware vCenter server based in cloud in either a normal Data Center or a Data Center of type Cloud, created with the appropriate cloud credentials.

To add VMware virtualization servers

1 Prerequisites:

See [“Prerequisites for adding VMware virtualization servers”](#) on page 169.

2 Navigate



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

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

3 Expand the data center > **Virtualization and Private Cloud** > **VMware** tab.

4 Launch the **+ vCenter** wizard.

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

- Specify the fully-qualified hostname of the vCenter Server that you want to discover along with its port number. The default port is 443. Support for IPv6 address is enabled. You can configure vCenter server using IPv6 address or hostname registered with DNS server.
- When entering login credentials, an administrative vCenter Server user account is required.
- If the data center has more than one IMS, a list of IMS names is shown. Select the IMS that you want to use to discover and monitor the vCenter Server and ESX servers.

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

If you choose the auto discover option, all currently available ESX servers are discovered. In addition, ESX servers later added to the vCenter Server are automatically discovered. However if some of the ESX servers associated with the vCenter Server are disconnected, then the auto discovery option is disabled.

If you are adding a vCenter server based in cloud, you have to select the auto discovery option as the ESX cluster membership changes are frequent and automatically initiated by the cloud provider.

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

The wizard notifies you of any issues.

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

For a vCenter server configured, if any of following changes are made in vCenter server after adding it to Resiliency Platform then,

- Any ESX server that is configured in Resiliency Platform is removed and re-added into vCenter server.
- Any existing ESX/Datastore cluster is destroyed and recreated.

After such a change in vCenter server, user should edit vCenter configuration in Resiliency Platform without any changes to the ESXs that were configured earlier. This allows the VMware NRT daemon in Resiliency Platform to start monitoring the objects that were re-added/re-created in vCenter server.

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

Prerequisites for adding VMware virtualization servers

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

- Ensure that the IMS can ping the vCenter servers or the ESX servers from which it can discover the information on VMware Infrastructure.
- If you want to have secure communication with the VMware vCenter server, ensure that you have installed the root CA certificate in the Resiliency Platform. Refer to
- Ensure that the vCenter Server user account that is used to add the servers to Resiliency Platform has the following privileges assigned on the data center level, depending upon your preferred data availability options:
 - See [“VMware vCenter Server privileges: configuring ESX servers”](#) on page 170.
 - See [“VMware vCenter Server privileges: basic monitoring of assets”](#) on page 171.
 - See [“VMware vCenter Server privileges: using third party replication”](#) on page 171.
 - See [“VMware vCenter Server privileges: replication using Resiliency Platform Data Mover in Hypervisor mode”](#) on page 173.
 - See [“VMware vCenter Server privileges: replication to cloud data center”](#) on page 177.
 -
 - See [“VMware vCenter Server privileges: using NetBackup”](#) on page 179.
 - See [“VMware vCenter Server privileges: for recovery of physical machines to VMware”](#) on page 181.

- You can configure a VMware vCenter to:
 - Start and stop the virtual machines.
 - Recover virtual machines to and from a remote premise based VMware data center
 - Recover virtual machines to and from a remote cloud based VMware data center
 - Recover virtual machines to a remote cloud data center
- For recovery of virtual machines in remote data center, Resiliency Platform supports following data availability choices:
 - Replication using Resiliency Platform Datamover using VAI/O interfaces (Hypervisor mode)
 - Replication using Resiliency Platform Datamover using an in-guest component (In-guest mode)
 - Replication using third party array replication technologies
 - Replication of backups of virtual machines using NetBackup AIR technology
 - Replication of backups of virtual machines to cloud using NetBackup Cloud Catalyst

VMware vCenter Server privileges: configuring ESX servers

To configure the ESX servers in Veritas Resiliency Platform, the vCenter Server user needs to have a role assigned that has the privileges, listed in the table, on the vCenter server data center level.

Review the following considerations before assigning the privileges:

- You need to maintain propagation of these data center level privileges on the child objects.
- The child objects should not restrict the data center level privileges.

Table 1-54 VMware vCenter Server privileges required for configuration of ESX servers and creation of resiliency groups

Privilege	Navigation path in VMware vCenter Server
Datstore.Browse	Datstore > Browse datastore
Host.Config.Network	Host > Configuration > Network configuration
StorageProfile.View	Profile-driven storage > Profile-driven storage view

Table 1-54 VMware vCenter Server privileges required for configuration of ESX servers and creation of resiliency groups (*continued*)

Privilege	Navigation path in VMware vCenter Server
System.Anonymous	These privileges need not be set explicitly.
System.Read	
System.View	

Before performing any operation on the resiliency groups, the privileges need to be elevated to match the privileges specified for performing operations.

See [“Prerequisites for adding VMware virtualization servers”](#) on page 169.

VMware vCenter Server privileges: basic monitoring of assets

To monitor (start and stop) the virtual machines in Veritas Resiliency Platform, the vCenter Server user needs to have a role assigned that has the privileges, listed in the table, on the vCenter server data center level.

Review the following considerations before assigning the privileges:

- You need to maintain propagation of these data center level privileges on the child objects.
- The child objects should not restrict the data center level privileges.

Table 1-55 VMware vCenter Server privileges required for basic monitoring of assets

Privilege	Navigation path in VMware vCenter Server
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off

See [“Prerequisites for adding VMware virtualization servers”](#) on page 169.

VMware vCenter Server privileges: using third party replication

To use any third party replication with VMware vCenter Server, the vCenter Server user needs to have a role assigned that has the privileges, listed in the table, on the vCenter server data center level.

Review the following considerations before assigning the privileges:

- You need to maintain propagation of these data center level privileges on the child objects.
- The child objects should not restrict the data center level privileges.

Table 1-56 VMware vCenter Server privileges required for disaster recovery using third party replication

Privilege	Navigation path in VMware vCenter Server
Datastore.AllocateSpace	Datastore > Allocate space
Datastore.Browse	Datastore > Browse datastore
Datastore.FileManagement	Datastore > Low level file operations
Datastore.Rename	Datastore > Rename
Datastore.Move	Datastore > Move
Host.Config.Network	Host > Configuration > Network configuration
Host.Config.Storage	Host > Configuration > Storage partition configuration
Network.Assign	Network > Assign network
Resource.ApplyRecommendation	Resource > Apply recommendation
Resource.AssignVMToPool	Resource > Assign virtual machine to resource pool
System.Anonymous	These privileges need not be set explicitly.
System.Read	
System.View	
VirtualMachine.Config.AddExistingDisk	Virtual machine > Configuration > Add existing disk
VirtualMachine.Config.AddNewDisk	Virtual machine > Configuration > Add new disk
VirtualMachine.Config.AddRemoveDevice	Virtual machine > Configuration > Add or remove device
VirtualMachine.Config.AdvancedConfig	Virtual machine > Configuration > Advanced
VirtualMachine.Config.Resource	Virtual machine > Configuration > Change resource
VirtualMachine.Config.EditDevice	Virtual machine > Configuration > Modify device settings
VirtualMachine.Config.RawDevice	Virtual machine > Configuration > Raw device
VirtualMachine.Config.RemoveDisk	Virtual machine > Configuration > Remove Disk
VirtualMachine.Config.Rename	Virtual machine > Configuration > Rename
VirtualMachine.Config.Settings	Virtual machine > Configuration > Settings

Table 1-56 VMware vCenter Server privileges required for disaster recovery using third party replication (*continued*)

Privilege	Navigation path in VMware vCenter Server
VirtualMachine.GuestOperations.Modify	Virtual machine > Guest operations > Guest operation modifications
VirtualMachine.GuestOperations.Execute	Virtual machine > Guest operations > Guest operation program execution
VirtualMachine.GuestOperations.Query	Virtual machine > Guest operations > Guest operation queries
VirtualMachine.Interact.DeviceConnection	Virtual machine > Interaction > Device connection
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on
VirtualMachine.Inventory.Register	Virtual machine > Inventory > Register
VirtualMachine.Inventory.Unregister	Virtual machine > Inventory > Unregister
VirtualMachine.Provisioning.Customize	Virtual machine > Provisioning > Customize
VApp.InstanceConfig	vApp > vApp instance configuration
VApp.ApplicationConfig	vApp > vApp application configuration

In addition to the above privileges, you also need the privileges required for configuring ESX servers:

See [“VMware vCenter Server privileges: configuring ESX servers”](#) on page 170.

See [“Prerequisites for adding VMware virtualization servers”](#) on page 169.

VMware vCenter Server privileges: replication using Resiliency Platform Data Mover in Hypervisor mode

To use Veritas Resiliency Platform Data Mover with VMware vCenter Server for replication to on-premises data center, the vCenter Server user needs to have a role assigned that has the privileges listed in the table, on the vCenter server data center level.

Review the following considerations before assigning the privileges:

- You need to maintain propagation of these data center level privileges on the child objects.
- The child objects should not restrict the data center level privileges.

Table 1-57 VMware vCenter Server privileges required for performing operations on veritas replication VIB

Operation	Privilege	Navigation path in VMware vCenter Server
Installation	Host.Cim.CimInteraction	Host > CIM > CIM interaction
	Host.Config.Maintenance	Host > Configuration > Maintenance
	Host.Config.Patch	Host > Configuration > Query patch
	Profile.Create	Host profile > Create
	StorageProfile.Update	Profile-driven storage > Profile-driven storage update
	StorageProfile.View	Profile-driven storage > Profile-driven storage view
Resolve installation errors	Host.Cim.CimInteraction	Host > CIM > CIM interaction
	Host.Config.Maintenance	Host > Configuration > Maintenance
Upgrade	Host.Cim.CimInteraction	Host > CIM > CIM interaction
	Host.Config.Maintenance	Host > Configuration > Maintenance
	Host.Config.Patch	Host > Configuration > Query patch
	Profile.Create	Host profile > Create
	StorageProfile.Update	Profile-driven storage > Profile-driven storage update
	StorageProfile.View	Profile-driven storage > Profile-driven storage view
Uninstallation	Host.Config.Patch	Host > Configuration > Query patch
	Host.Config.Maintenance	Host > Configuration > Maintenance

Table 1-58 VMware vCenter Server privileges required for using Resiliency Platform Data Mover for replication using Resiliency Platform Data Mover in Hypervisor mode

Privilege	Navigation path in VMware vCenter Server
Datastore.AllocateSpace	Datastore > Allocate space
Datastore.FileManagement	Datastore > Browse datastore
Datastore.Browse	Datastore > Low level file operations
Host.Cim.CimInteraction	Host > CIM > CIM interaction

Table 1-58 VMware vCenter Server privileges required for using Resiliency Platform Data Mover for replication using Resiliency Platform Data Mover in Hypervisor mode (*continued*)

Privilege	Navigation path in VMware vCenter Server
Host.Config.Network	Host > Configuration > Network configuration
Host.Config.Storage	Host > Configuration > Storage partition configuration
Host.Config.Maintenance	Host > Configuration > Maintenance
Host.Config.NetService	Host > Configuration > Security profile and firewall
Profile.Create	Host profile > Create
Network.Assign	Network > Assign network
StorageProfile.Update	Profile-driven storage > Profile-driven storage update
StorageProfile.View	Profile-driven storage > Profile-driven storage view
StoragePod.Config	DatastoreCluster > Configure Datastore Cluster
Resource.ApplyRecommendation	Resource > Apply recommendation
Resource.AssignVMToPool	Resource > Assign virtual machine to resource pool
System.Anonymous	These privileges need not be set explicitly.
System.Read	
System.View	

Table 1-58 VMware vCenter Server privileges required for using Resiliency Platform Data Mover for replication using Resiliency Platform Data Mover in Hypervisor mode (*continued*)

Privilege	Navigation path in VMware vCenter Server
VirtualMachine.Config.AddExistingDisk	Virtual machine > Configuration > Add existing disk
VirtualMachine.Config.AddNewDisk	Virtual machine > Configuration > Add new disk
VirtualMachine.Config.AddRemoveDevice	Virtual machine > Configuration > Add or remove device
VirtualMachine.Config.AdvancedConfig	Virtual machine > Configuration > Advanced
VirtualMachine.Config.Resource	Virtual machine > Configuration > Change resource
VirtualMachine.Config.EditDevice	Virtual machine > Configuration > Modify device settings
VirtualMachine.Config.RawDevice	Virtual machine > Configuration > Raw device
VirtualMachine.Config.RemoveDisk	Virtual machine > Configuration > Remove disk
VirtualMachine.Config.Rename	Virtual machine > Configuration > Rename
VirtualMachine.Config.Settings	Virtual machine > Configuration > Settings
VirtualMachine.Config.ExtendDisk	Virtual machine > Configuration > Extend Virtual Disk
VirtualMachine.GuestOperations.Modify	Virtual machine > Guest operations > Guest operation modifications
VirtualMachine.GuestOperations.Execute	Virtual machine > Guest operations > Guest operation program execution
VirtualMachine.GuestOperations.Query	Virtual machine > Guest operations > Guest operation queries
VirtualMachine.Interact.DeviceConnection	Virtual machine > Interaction > Device connection
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on
Virtual machine.Interaction.Configure CD media	Virtual Machine > Interaction > Configure CD
VirtualMachine.Inventory.CreateFromExisting	Virtual machine > Inventory > Create from existing
VirtualMachine.Inventory.Create	Virtual machine > Inventory > Create new
VirtualMachine.Inventory.Register	Virtual machine > Inventory > Register
VirtualMachine.Inventory.Delete	Virtual machine > Inventory > Remove
VirtualMachine.Inventory.Unregister	Virtual machine > Inventory > Unregister

Table 1-58 VMware vCenter Server privileges required for using Resiliency Platform Data Mover for replication using Resiliency Platform Data Mover in Hypervisor mode (*continued*)

Privilege	Navigation path in VMware vCenter Server
VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.Customize	Virtual machine > Provisioning > Clone virtual machine Virtual machine > Provisioning > Customize
VirtualMachine.State.CreateSnapshot VirtualMachine.State.RemoveSnapshot	Virtual machine > Snapshot management > Create snapshot Virtual machine > Snapshot management > Remove snapshot
VApp.InstanceConfig VApp.ApplicationConfig	vApp > vApp instance configuration vApp > vApp application configuration

In addition to the above privileges, you also need the privileges required for configuring ESX servers:

See [“VMware vCenter Server privileges: configuring ESX servers”](#) on page 170.

See [“Prerequisites for adding VMware virtualization servers”](#) on page 169.

VMware vCenter Server privileges: replication to cloud data center

To use Veritas Resiliency Platform Data Mover with VMware vCenter Server for replication to any cloud data center, the vCenter Server user needs to have a role assigned that has the privileges, listed in the table, on the vCenter server data center level.

Review the following considerations before assigning the privileges:

- You need to maintain propagation of these data center level privileges on the child objects.
- The child objects should not restrict the data center level privileges.

Table 1-59 VMware vCenter Server privileges required for using Resiliency Platform Data Mover for replication to cloud data center

Privilege	Navigation path in VMware vCenter Server
Datastore.AllocateSpace	Datastore > Allocate space
Datastore.FileManagement	Datastore > Low level file operations
Datastore.Browse	Datastore > Browse datastore

Table 1-59 VMware vCenter Server privileges required for using Resiliency Platform Data Mover for replication to cloud data center
(continued)

Privilege	Navigation path in VMware vCenter Server
Host.Config.Network	Host > Configuration > Network configuration
Host.Config.Storage	Host > Configuration > Storage partition configuration
Network.Assign	Network > Assign network
Resource.ApplyRecommendation	Resource > Apply recommendation
Resource.AssignVMToPool	Resource > Assign virtual machine to resource pool
System.View	These privileges need not be set explicitly.
System.Anonymous	
System.Read	
VirtualMachine.Config.AddExistingDisk	Virtual machine > Configuration > Add existing disk
VirtualMachine.Config.AddNewDisk	Virtual machine > Configuration > Add new disk
VirtualMachine.Config.AddRemoveDevice	Virtual machine > Configuration > Add or remove device
VirtualMachine.Config.AdvancedConfig	Virtual machine > Configuration > Advanced
VirtualMachine.Config.Resource	Virtual machine > Configuration > Change resource
VirtualMachine.Config.DiskExtend	Virtual machine > Configuration > Extend virtual disk
VirtualMachine.Config.EditDevice	Virtual machine > Configuration > Modify device settings
VirtualMachine.Config.RawDevice	Virtual machine > Configuration > Raw device
VirtualMachine.Config.RemoveDisk	Virtual machine > Configuration > Remove disk
VirtualMachine.Config.Rename	Virtual machine > Configuration > Rename
VirtualMachine.Config.Settings	Virtual machine > Configuration > Settings
VirtualMachine.Interact.DeviceConnection	Virtual machine > Interaction > Device connection
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on

Table 1-59 VMware vCenter Server privileges required for using Resiliency Platform Data Mover for replication to cloud data center
(continued)

Privilege	Navigation path in VMware vCenter Server
VirtualMachine.Inventory.CreateFromExisting	Virtual machine > Inventory > Create from existing
VirtualMachine.Inventory.Create	Virtual machine > Inventory > Create new
VirtualMachine.Inventory.Register	Virtual machine > Inventory > Register
VirtualMachine.Inventory.Delete	Virtual machine > Inventory > Remove
VirtualMachine.Inventory.Unregister	Virtual machine > Inventory > Unregister
VirtualMachine.Provisioning.Clone	Virtual machine > Provisioning > Clone virtual machine
VirtualMachine.Provisioning.Customize	Virtual machine > Provisioning > Customize
VirtualMachine.State.CreateSnapshot	Virtual machine > Snapshot management > Create snapshot
VirtualMachine.State.RemoveSnapshot	Virtual machine > Snapshot management > Remove snapshot

In addition to the above privileges, you also need the privileges required for configuring ESX servers:

See [“VMware vCenter Server privileges: configuring ESX servers”](#) on page 170.

See [“Prerequisites for adding VMware virtualization servers”](#) on page 169.

VMware vCenter Server privileges: using NetBackup

To recover VMware virtual machine from NetBackup generated backup images to the recovery data center, the vCenter Server user needs to have a role assigned that has the privileges, listed in the table, on the vCenter server data center level.

Review the following considerations before assigning the privileges:

- You need to maintain propagation of these data center level privileges on the child objects.
- The child objects should not restrict the data center level privileges.

Table 1-60 VMware vCenter Server privileges required for using NetBackup integration

Privilege	Navigation path in VMware vCenter Server
Datastore.AllocateSpace	Datastore > Allocate space
Datastore.FileManagement	Datastore > Low level file operations
Datastore.Browse	Datastore > Browse datastore
Host.Config.Network	Host > Configuration > Network configuration
Network.Assign	Network > Assign network
Resource.ApplyRecommendation	Resource > Apply recommendation
Resource.AssignVMToPool	Resource > Assign virtual machine to resource pool
System.View System.Anonymous System.Read	These privileges need not be set explicitly.
VirtualMachine.Config.AddRemoveDevice	Virtual machine > Configuration > Add or remove device
VirtualMachine.Config.AdvancedConfig	Virtual machine > Configuration > Advanced
VirtualMachine.Config.Resource	Virtual machine > Configuration > Change resource
VirtualMachine.Config.EditDevice	Virtual machine > Configuration > Modify device settings
VirtualMachine.Config.Settings	Virtual machine > Configuration > Settings
VirtualMachine.Config.DiskExtend	Virtual machine > Configuration > Extend virtual disk
VirtualMachine.GuestOperations.Modify	Virtual machine > Guest operations > Guest operation modifications
VirtualMachine.GuestOperations.Execute	Virtual machine > Guest operations > Guest operation program execution
VirtualMachine.GuestOperations.Query	Virtual machine > Guest operations > Guest operation queries

Table 1-60 VMware vCenter Server privileges required for using NetBackup integration (*continued*)

Privilege	Navigation path in VMware vCenter Server
VirtualMachine.Interact.DeviceConnection	Virtual machine > Interaction > Device connection
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on
VirtualMachine.Inventory.Create	Virtual machine > Inventory > Create new
VirtualMachine.Inventory.Register	Virtual machine > Inventory > Register
VirtualMachine.Inventory.Delete	Virtual machine > Inventory > Remove
VirtualMachine.Inventory.Unregister	Virtual machine > Inventory > Unregister
VirtualMachine.Provisioning.Customize	Virtual machine > Provisioning > Customize

In addition to the above privileges, you also need the privileges required for configuring ESX servers:

See [“VMware vCenter Server privileges: configuring ESX servers”](#) on page 170.

See [“Prerequisites for adding VMware virtualization servers”](#) on page 169.

VMware vCenter Server privileges: for recovery of physical machines to VMware

To use Veritas Resiliency Platform Data Mover with VMware vCenter Server for recovery to physical machines to VMware, the vCenter Server user needs to have a role assigned that has the privileges, listed in the table, on the vCenter server data center level.

Review the following considerations before assigning the privileges:

- You need to maintain propagation of these data center level privileges on the child objects.
- The child objects should not restrict the data center level privileges.

Table 1-61 VMware vCenter Server privileges required for recovery of physical machines to VMware

Privileges	Navigation path in VMware vCenter Server
Datastore.AllocateSpace	Datastore > Allocate space
Datastore.Browse	Datastore > Browse datastore
Datastore.FileManagement	Datastore > Low level file operations
Host.Config.Network	Host > Configuration > Network configuration
Host.Config.SystemManagement	Host > Configuration > SystemManagement
Network.Assign	Network > Assign network
Resource.ApplyRecommendation	Resource > Apply recommendation
Resource.AssignVMToPool	Resource > Assign virtual machine to resource pool
StoragePod.Config	StoragePod > Config
StorageProfile.View	Profile-driven storage > Profile-driven storage view
System.Anonymous	These privileges need not be set explicitly.
System.Read	
System.View	
VApp.ApplicationConfig	vApp > vApp application configuration
VApp.InstanceConfig	vApp > vApp instance configuration

Table 1-61 VMware vCenter Server privileges required for recovery of physical machines to VMware (*continued*)

Privileges	Navigation path in VMware vCenter Server
VirtualMachine.Config.AddExistingDisk	Virtual machine > Configuration > Add existing disk
VirtualMachine.Config.AddNewDisk	Virtual machine > Configuration > Add new disk
VirtualMachine.Config.AddRemoveDevice	Virtual machine > Configuration > Add or remove device
VirtualMachine.Config.AdvancedConfig	Virtual machine > Configuration > Advanced
VirtualMachine.Config.DiskExtend	Virtual machine > Configuration > Extend virtual disk
VirtualMachine.Config.EditDevice	Virtual machine > Configuration > Modify device settings
VirtualMachine.Config.RawDevice	Virtual machine > Configuration > Raw device
VirtualMachine.Config.RemoveDisk	Virtual machine > Configuration > Remove disk
VirtualMachine.Config.Rename	Virtual machine > Configuration > Rename
VirtualMachine.Config.Resource	Virtual machine > Configuration > Change resource
VirtualMachine.Config.Settings	Virtual machine > Configuration > Settings
VirtualMachine.Interact.Configure CD Media	Virtual Machine > Interaction > Configure CD
VirtualMachine.Interact.DeviceConnection	Virtual machine > Interaction > Device connection
VirtualMachine.Interact.PowerOff	Virtual machine > Interaction > Power off
VirtualMachine.Interact.PowerOn	Virtual machine > Interaction > Power on
VirtualMachine.Interact.VMware Tools Install	Virtual machine > Interaction > VMware Tools Install
VirtualMachine.Inventory.CreateFromExisting	Virtual machine > Inventory > Create from existing
VirtualMachine.Inventory.Create	Virtual machine > Inventory > Create new
VirtualMachine.Inventory.Delete	Virtual machine > Inventory > Remove
VirtualMachine.Inventory.Register	Virtual machine > Inventory > Register
VirtualMachine.Inventory.Unregister	Virtual machine > Inventory > Unregister

Table 1-61 VMware vCenter Server privileges required for recovery of physical machines to VMware (*continued*)

Privileges	Navigation path in VMware vCenter Server
VirtualMachine.Provisioning.Clone VirtualMachine.Provisioning.Customize	Virtual machine > Provisioning > Clone virtual machine Virtual machine > Provisioning > Customize
VirtualMachine.State.CreateSnapshot VirtualMachine.State.RemoveSnapshot	Virtual machine > Snapshot management > Create snapshot Virtual machine > Snapshot management > Remove snapshot

In addition to the above privileges, you also need the privileges required for configuring ESX servers:

See [“VMware vCenter Server privileges: configuring ESX servers”](#) on page 170.

See [“Prerequisites for adding VMware virtualization servers”](#) on page 169.

Preparing host for replication

To enable the replication in Resiliency Platform using Resiliency Platform Data Mover, you need to add the asset and prepare it for replication. Asset can be a physical machine or a virtual machine.

To prepare a host for replication

- 1 Ensure that you understand the use cases and prerequisites for adding hosts to an IMS.
 - See [“About adding host assets”](#) on page 503.
See [“Prerequisites for adding hosts”](#) on page 505.

Note: After you add a new data disk to a Windows host and attach it to an IDE controller, you need to initialize the disk. This needs to be done before performing the Prepare Host for Replication task.

- 2 Navigate to **Settings** (menu bar) > **Infrastructure** > **Details View**.
You can also access this page from the **Quick Actions** menu.
- 3 Go to the on-premises data center and click **Data Mover**.
- 4 Under **Resiliency Platform Data Mover**, click **Prepare host for replication**.

- 5 In the wizard, select the type of host to be added.
 - Discovered Virtual Machines: Choose this option to select one or more virtual machines that are already discovered from your virtualization environment.
 - Non-discovered hosts (no-hypervisor configured or physical hosts): Choose this option to add virtual machines which are not yet discovered or which are physical hosts. You can import multiple hosts from a simple comma separated text file using this option.
- 6 Choose the desired configuration mode from the following options:
 - Using VRP Console: You can configure a host using the username and password of the host. Resiliency Platform automatically installs the appropriate host agent packages on the host.
 - Using pre-deployed VRP Host Agent Packages: Choose this option if you do not have the username and password of the host or if you require full control for deploying the host agent packages using your tools and then configure the host without providing the username and password from the Resiliency Platform console. The VRP host agent packages can be downloaded from the IMS.
See [“Downloading Resiliency Platform host agent packages from IMS”](#) on page 501.
See [“Configuring Resiliency Platform host agent packages manually”](#) on page 501.

Note: Ensure that you choose the same IMS to add this host from which you chose to download the VRP host agent package bundle to install and configure on that host.

- 7 Select the IMS to which you want to add a host and click **Next**
- 8 Based on the selected Host Type and Configuration Mode specify the following inputs:

Host Type	Configuration mode	Required Inputs
Discovered Virtual Machines	Using VRP console	<ul style="list-style-type: none"> ■ Select one or more Virtual Machines and click Next. You can filter the virtual machines based on environment or on the virtualization server they are running on, or search them using their virtual machine name, family type, or host name. ■ Review the discovered Hostname for the virtual machine and modify it, if necessary. ■ Enter the username and password for each virtual machine in the table. See Prerequisites for adding host for information about using non- root user accounts to add Linux hosts. ■ Click  to delete a row. <p>See “Prerequisites for adding hosts ” on page 505. for information about using non-root user accounts to add Linux hosts.</p>
Discovered Virtual Machines	Using a pre-deployed VRP Host Agent	<ul style="list-style-type: none"> ■ Select one or more Virtual Machines and click Next. You can filter the virtual machines based on environment or on the virtualization server they are running on, or search them using their virtual machine name, family type, or host name. ■ Review the discovered Hostname for the virtual machine and modify it, if necessary.

Host Type	Configuration mode	Required Inputs
Non-discovered hosts Physical hosts	Using VRP console	<p>Type the host name and user credentials information in the table row. Use the following options to add information for multiple hosts.</p> <ul style="list-style-type: none"> ■ Use the <i>Add new row</i> option to add a blank table row. ■ Use the <i>Import from a text file</i> option to import information from a text file with comma separated values. To do this, select the desired text file and click Load host details. ■ Use the  icon to copy the details of the elected table row. You can edit the details of the newly added row. ■ Use the  icon to delete the desired row. <p>See “Prerequisites for adding hosts ” on page 505. for information about using non-root user accounts to add Linux hosts.</p>

Host Type	Configuration mode	Required Inputs
Non-discovered hosts or Physical hosts	Using a pre-deployed VRP Host Agent	<p>Type the host name in the table row. Use the following options to add information for multiple hosts.</p> <ul style="list-style-type: none"> ■ Use the <i>Add new row</i> option to add a blank table row. ■ Use the <i>Import from a text file</i> option to import information from a text file with comma separated values. To do this, select the desired text file and click Load host details. ■ Use the  icon to copy the details of the selected table row. You can edit the details of the newly added row. ■ Use the  icon to delete the desired row.

- 9 Once you are done with entering the data for all the hosts, click **Submit** to initiate the Prepare Host for Replication activity for each host.
- 10 Click the **Status** column value to navigate to the respective activity details view for a host to check if the activity is in progress or has failed.

The host is listed in the table for Replication Hosts on the Resiliency Platform Data Mover tab with Status column as Connected when the activity completes successfully.

When the recovery environment is AWS Cloud, the cloud storage and network drivers must be installed on the host in following conditions:

- Windows Paravirtual (PV) drivers to recover Windows hosts
- Xen block storage and networking drivers to recover SUSE Linux 11.4 hosts

If the compatible drivers are not installed already, the bundled versions of the required drivers are installed automatically when the hosts are added to resiliency group.

The bundled versions of the driver are located at:

- For SUSE Linux platform: `/var/opt/VRTSsfmh/spool/addons/store/
<VRTSitrptapversion_of_the_release>/AWSPVDriver`

Refer to the `Readme.txt` file in this directory for version information of the bundled drivers.

Managing multiple VMWare Virtual machines with same BIOS ID

If a VMWare virtual machine is already protected with Veritas Standalone Replication and has in-guest component then you cannot provision another virtual machine with the same BIOS ID as that of the protected virtual machine unless the protected virtual machine is upgraded to version 3.5.

Downloading Resiliency Platform host agent packages from IMS

You can download the Resiliency Platform host agent packages compressed bundle from IMS to which you intend to add the hosts. When you deploy and configure the downloaded Resiliency Platform host agent package on a host, you can prepare it for replication from the Resiliency Platform console without user name and password. To add the host in the Prepare Host for Replication wizard, you have to select the same IMS from which you downloaded the host agent package.

To download the Resiliency Platform host agent packages from an IMS

- 1 Navigate to **Settings** (menu bar) > **Infrastructure** > **Details View**.
You can also access this page from the **Quick Actions** menu.
- 2 Go to the on-premises data center and locate the desired IMS card.
- 3 Click **Vertical Ellipsis** menu on the IMS card and click **Resiliency Platform Host Agent Package**.
- 4 The popup displays the host agent download URL. Use the Copy button to copy the link to the clipboard.
Alternatively, you may manually select the link and copy it.
- 5 Browse to the URL and download the zip bundle.

Configuring Resiliency Platform host agent packages manually

You can manually install and configure Resiliency Platform host agent packages on a host before preparing it for replication using Resiliency Platform Data Mover. The host agent package zip bundle can be downloaded from an IMS.

The Resiliency Platform host agent bundle contains a configuration script which configures the host so that the specific IMS can communicate with the host without using the host credentials.

Before you begin configuring the Resiliency Platform host agent packages, download the Resiliency Platform host agent packages zip bundle from the IMS on the host and extract it at a temporary location, for example, `/tmp` on Linux host or `c:\temp` on Windows host.

See [“Downloading Resiliency Platform host agent packages from IMS”](#) on page 501.

To configure Resiliency Platform host agent packages manually on a host

- 1 Install the packages on required host.
 - On Linux host
 - Separate folders for all supported platforms are pre-created inside Linux folder. Go to folder depending on which platform the package is to be installed. Copy `VRTSitrptap-<version>.rpm` to a temporary folder (for example: `/tmp`) on Linux host.
 - Install `VRTSsfmh-<version>.rpm` on Linux host.
 - Install `VRTSitrptap-<version>.rpm` on Linux host.
 - On Windows host
 - Copy `VRTSitrptap-<version>.exe` from `x86_64` folder on Windows host.
 - Install `VRTSsfmh-<version>.msi` on Windows host.
 - Install `VRTSitrptap-<version>.exe` on Windows host.
 - Restart the Windows host so that the replication driver is loaded.
- 2 Run the bundled script to configure the host for communication from IMS.
 - On Linux host
 - Execute the `ConfigureHostFor-<IMS-Hostname/IP>.pl` script on Linux host
as:

```
# /opt/VRTSsfmh/bin/perl /tmp/ConfigureHostFor-<IMS-Hostname/IP>.pl --configure
```
 - On Windows host
 - Execute the `ConfigureHostFor-<IMS-Hostname/IP>.pl` script on Windows host as:

```
C:\> "C:\Program Files\Veritas\VRTSsfmh\bin\perl.exe" C:\temp\ConfigureHostFor-<IMS-Hostname/IP>.pl --configure
```

Configuring password-less sudo privileges for user account on Linux host

You can add a Linux host to an IMS or move it to another IMS using a user account that does not belong to root group, provided that the user account has sudo privileges to run the following commands without requiring a sudo password:

- `/bin/rpm`
- `/opt/VRTSsfmh/bin/perl`

Refer to the sample sudo configuration entry below for a user account named `user1`:

```
user1 ALL=(root) NOPASSWD: /bin/rpm,/opt/VRTSsfmh/bin/perl
```

Note: Ensure that the sudo binary path is included in the system PATH environment.

If the user account is part of a group for which sudo configuration is defined, the corresponding sudo configuration section for that group should have an entry for that user account to run commands with NOPASSWD label.

Versions like Linux RHEL 6.8, RHEL 7.2 and Centos 6.7, by default enable the `requiretty` property in the sudo configuration. This property must be disabled to allow running commands using sudo without a password prompt. To disable the property, comment the line `Defaults requiretty` in the sudo configuration. A sudo user (in case of linux workloads) requires the sudo package version to be minimum 1.7.8.

About adding host assets

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

Table 1-62 Use cases for adding host assets

Use case	Details
Application discovery and management	<p>For discovery of supported applications on either physical systems or virtual machines, you must add the physical system or virtual machine as a host.</p> <p>Note: For the use case of discovering and managing virtual machines rather than applications, you must add the virtual machines as hosts.</p> <p>For discovery of a custom application, after you add the hosts, you must also add the application on the Assets page.</p>
VMware vCenter Server discovery (optional)	<p>You can add a host to be used by the IMS for discovery of a VMware vCenter Server.</p>
Hardware replication	<p>For storage array-based replication, you may need to install array-specific software on a host and add the host as a discovery host.</p> <p>More information is available on requirements for adding enclosures for array-based replication.</p>
Resiliency Platform Data Mover host	<p>To manage and protect virtual machines using Resiliency Platform Data Mover, you need to add the virtual machines as hosts.</p>
Manage physical machines	<p>To manage and protect physical machines, you need to add the physical machines as hosts.</p>

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

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

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

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

See [“Prerequisites for adding hosts”](#) on page 505.

Prerequisites for adding hosts

Before you add hosts to Veritas Resiliency Platform for discovery and monitoring by an Infrastructure Management Server (IMS), ensure that the following prerequisites are fulfilled.

General prerequisites for adding host assets:

- Ensure that the IMS can communicate with the host.
- Ensure that the time difference between the system clocks on the IMS and on the host is not more than 90 minutes. It is recommended to configure NTP on the virtual machine that needs to be secured. NTP should be configured in such a way that when the virtual machine is on the source data center, it remains in sync with the IMS on source data center. After migration, the virtual machine should be in sync with the IMS on the target data center.
- If a text file with comma separated values is used to add hosts, ensure that it uses the correct syntax.
- Ensure that the vCenter server and the hosts which are being prepared for in-guest protection need to be added to the same IMS and that they cannot be part of two different IMS.

Additional prerequisites for Linux systems

- To add a Linux host with the configuration mode option, *Configure using VRP Console*, use a user account that is a part of the root group or a user account that has password-less sudo privileges.
See [“Configuring password-less sudo privileges for user account on Linux host”](#) on page 503.
- Ensure that `openssh-clients` package is present in the system. Typically it is installed in the operating system by default.
- For Resiliency Platform Data Mover host, it is recommended to install and configure `ntpd`. It ensures that the system’s time remains in sync even after migration to the other site.
- For Resiliency Platform Data Mover host, ensure that `dmidecode` package is available in the system. Typically it is installed in the operating system by default. However, it may be not present in minimal OS installations.
- In order to install the host package while adding the Linux host, ensure that the `PasswordAuthentication` field is set to **yes** in the `/etc/ssh/sshd_config` file on the host.
- You need to remove all the stale network files from the below directory before adding the host.
 - **RHEL and Centos path:** `/etc/sysconfig/network-scripts`

- **SUSE:** `/etc/sysconfig/network`
- If a host is being added or prepared using a sudo user (in case of linux workloads), it requires the sudo package version to be minimum 1.7.8.
- For Suse 15, ensure that `insserv` package is installed in the system. Typically, it is installed in the operating system by default.
- In order to persist the iptable rules, ensure that the iptables service is running and the `IPTABLES_SAVE_ON_STOP`, `IPTABLES_SAVE_ON_RESTART` and `IPTABLES_SAVE_COUNTER` directives are set to **yes** in the iptables config file.

Additional prerequisites for Windows systems

- To add a Windows host, it is recommended to use a domain user account with local administrator privileges.
If you cannot use a domain user account with local administrator privileges, you have an option to use an Administrator user or a user in local administrator group with required prerequisites.
See [“User account required for adding a Windows host”](#) on page 507.
- Ensure that you have appropriate User Access Control (UAC) settings for the user that is used for adding the Windows host.
See [“UAC settings required for adding a Windows host”](#) on page 507.
- The Windows Management Instrumentation (WMI) service must be running.
- Ensure that the password for the host does not contain a double quotes character.
- After you add a new data disk to a Windows host and attach it to IDE controller, you need to initialize the disk. This needs to be done before adding the host.
- If you have McAfee antivirus already installed on the virtual machines, you need to disable the other encryption options before adding the host.

Additional prerequisites for Oracle Discovery Hosts

- VRP supports Redhat Enterprise Linux hosts as Oracle Discovery Hosts.
- You have to install Oracle Client (Net) on the host. Oracle Net on the host must be configured by “oracle” user such that the `tnsnames.ora` points to all the Oracle instances to be discovered and managed. You can use the Oracle Net Configuration Assistance (`netca`) to configure the `tnsnames.ora` file.
- If you are configuring Oracle RAC, ensure that all RAC nodes are configured in the `tnsnames.ora` file. Make sure that `sqlplus` command on the discovery host is able to reach all the Oracle database instances individually.

- If the Oracle Database is replicated using Oracle DataGuard, make sure that you add only the database that is local to the Datacenter in the tnsnames.ora. You should use a different Discovery Host in the remote Datacenter for managing the remote Database.
- You may use the same discovery host to manage the multiple Oracle databases in the same datacenter. The Oracle Net has to be configured to be able to communicate with all the instances of all these databases.

See “[About adding host assets](#)” on page 503.

User account required for adding a Windows host

To add a Windows host, you need to have any one of the following credential sets:

Table 1-63 Credentials required for adding a Windows host

Windows host being added	Prerequisite
Domain user	Must have local administrator privileges on the Windows host being added.
Administrator	None
User in local administrators group	Add a registry entry to disable the remote restriction on the Windows host being added: Microsoft documentation Once the Windows host is added, you can enable the remote restriction again on that host.

UAC settings required for adding a Windows host

Following is the list of User Account Control (UAC) settings required for adding a Windows host in Veritas Resiliency Platform.

Table 1-64 UAC settings required for adding a Windows host

UAC Policy	Local administrator user	Domain user with administrator privileges	Domain user with local administrator privileges
Admin Approval Mode for Built-in Administrator Account	Disabled	Enabled	Enabled

Table 1-64 UAC settings required for adding a Windows host (*continued*)

UAC Policy	Local administrator user	Domain user with administrator privileges	Domain user with local administrator privileges
Allow UIAccess applications to prompt for elevation without using the secure desktop	Enabled	Disabled	Disabled
Behavior of the elevation prompt for administrator in Admin Approval Mode	Prompt for credentials on secure desktop	Prompt for credentials on secure desktop	Prompt for credentials on secure desktop
Behavior of the elevation prompt for administrator for standard users	Prompt for credentials	Prompt for credentials	Prompt for credentials
Detect application installations and prompt for elevation	Enabled	Enabled	Enabled
Only elevate executables that are signed and validated	Enabled	Enabled	Enabled
Only elevate UI Access application that are installed in secure locations	Enabled	Enabled	Enabled
Run all administrators in Admin Approval Mode	Enabled	Enabled	Enabled
Switch to the secure desktop when prompting for elevation	Disabled	Disabled	Disabled
Virtualize file and registry write failures to per-user locations	Enabled	Enabled	Enabled

Infrastructure Pairing

For recovering assets to the respective cloud platform, you have to do following infrastructure pairing. Refer the following topics:

-
- See [“About network objects”](#) on page 509.

- See [“Network pairs for recovering virtual machines to Google Cloud Platform \(GCP\)”](#) on page 512.

About network objects

Resiliency Platform discovers and displays information about layer 2 and layer 3 networks for the discovered assets.

Layer 2: The second level in the seven-layer OSI reference model, is used to transfer data between adjacent networks in a WAN or LAN environment. This layer is also known as Data Link Layer.

Examples of layer 2 networks: Port group/VLAN, vSwitch, cloud network and cloud subnet if the target data center is in cloud.

Layer 3: The network layer in the OSI reference model. mainly include routing and forwarding, as well as internetworking, addressing, packet sequencing, congestion control and further error handling.

Examples of layer 3 networks: Subnets and cloud subnets.

Resiliency Platform discovers and displays information about Layer 2 and Layer 3 networks in your datacenter. For VMware datacenter, Resiliency Platform 4.0 also supports NSX-T type networks present in vCenter server. It discovers and displays information about NSX-T provisioned L2 network objects from vCenter like Opaque switch, Opaque network and virtual distributed switch 7.0. NSX-T transport zones were discovered as vSwitches and NSX-T segments are discovered as port groups/VLAN.

Note: While adding transport zone in NSX-T Manager, either select NVDS or VDS for all the hosts.

For cloud technologies like AWS, Azure, vCloud Director and Google Cloud Platform, cloud subnets serve the purpose for both the layer 2 and layer 3 networks in the network pair. Network objects like private cloud subnet and private cloud network are listed under **Network Types** drop down irrespective of the cloud data center configuration.

You can manually add subnets, VMware port group/VLANs and Hyper-V VLANs that are not discovered in a data center. You cannot add vSwitches and cloud networks, if they are not discovered. Adding subnets using IPv6 address and pairing them across data centers is supported. Either while adding the network objects or while editing the discovered network objects, you need to choose a purpose. Purpose can be Production or Rehearsal.

See [“About Purpose”](#) on page 511.

For mapping the purpose of the network objects as Production or Rehearsal,

Network objects like private cloud subnet and private cloud network are listed under **Network Types** drop down irrespective of the cloud data center configuration.

A network pair is created using network objects across data centers. The network pairs should be defined before a resiliency group is created. Depending on whether the participating networks in the pair are layer 2 or layer 3 networks, the pair can be used for connecting the assets to the networks, or for assisting the customization of the IP addresses in the target network. The create network pair operation eliminates the need to manually connect each asset to a network at the target data center. For example, port group/VLAN to cloud network and subnet to port group/VLAN.

When the resiliency group is created, the network objects in the network pairs are evaluated by the Resiliency Platform. The CIDR (Classless Inter-Domain Routing) information from the network object is used to automatically calculate the IP address for the applicable assets in the respective target network if any of the below mapping is done:

- Subnet to subnet
- Subnet to cloud subnet
- Cloud subnet to cloud subnet
- Private Cloud Subnet to Private Cloud Subnet

The layer 2 network object pairs are mandatory to be defined for recovery to cloud environments, private cloud environments and recovery of physical machines to VMware environment. This mapping is optional for recovery from on-premises to on-premises environment. If the mappings are not defined for recovery to on-premise environment, then the virtual machine NICs are not connected to any network.

A layer 3 network pair is optional. If it is defined, the IP address for the asset is calculated based on the target subnet CIDR and can be further customized. If the network pair is not defined, then the IP address for the adapter gets assigned in one of the following ways:

1. If the IP customization option is checked, user must enter the IP address (IPv4 or IPv6 address depending upon the network configured) that needs to be assigned to the virtual machine NIC.
2. If the IP customization option is not checked, a DHCP (Dynamic Host Configuration Protocol) IP address gets assigned to the virtual machine adapter if the target technology supports it. (For example: cloud environments).
3. If the IP customization option is not checked, for on-premises to on-premises recovery, the virtual machine adapter IP settings are not changed.

When you perform a migrate, recover or rehearsal operation on a resiliency group, the Resiliency Platform evaluates the network pairs that have the layer 2 network objects and gets connected to the expected target network.

Using Resiliency Platform console, you can create network groups of cloud subnets for AWS cloud data center and port group/VLAN for VMware environment only.

In case of Google Cloud Platform, if shared subnets are discovered, the host project name is appended to the VPC name to distinguish the shared subnets from another project.

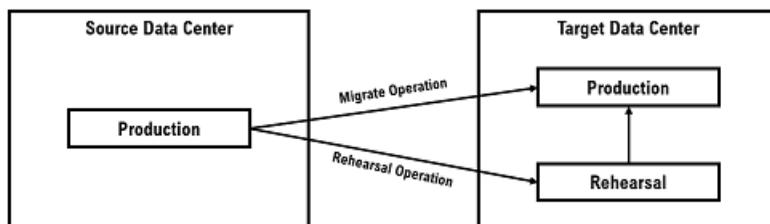
When you clone your virtual machines, ensure that you assign the appropriate host name and IP address to the cloned virtual machines.

About Purpose

When the corresponding network objects are involved in recovery tasks such as migrate, recover, or resync, then set the purpose as Production on those network objects. Similarly, when the corresponding network objects are involved only in Rehearsal operations, then set the purpose as Rehearsal on those network objects.

To perform recovery operations such as migrate, recover, or resync, you need to pair the network objects with a purpose as Production from source data center with network objects with a purpose as Production in the target data center. IPv4 network objects (for example, subnets) can be paired with only IPv4 network objects, and IPv6 network objects can be paired with only IPv6 network objects.

When you want to perform the rehearsal operations, you can map the network objects with purpose as Production with network objects (corresponding to rehearsal networks) with purpose as Rehearsal in the target data center. For example, you can map the AWS Cloud Subnet (Production as purpose) to AWS Cloud Subnet (Rehearsal as purpose). Following figure explains the mapping across the data center:



The create network pair operation eliminates the need to manually connect each virtual machine to a network at the recovery data center.

After you have paired the networks successfully, the target networks and the IP addresses are computed programmatically, and applied to the virtual machines.

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

For example:

In your source data center you have two subnets with Production and rehearsal as the purpose. The Production subnet is mapped with the Rehearsal subnet. A similar setup is present on your target data center. Both the subnets having the purpose as Production are paired.

Now when you perform the migrate operation, since the subnets having purpose as Production are paired, the virtual machine is migrated with appropriate network settings.

For example, a virtual machine has IP address 10.20.30.40 and is part of subnet 10.20.30.0/24 on source data center. This subnet is paired with another subnet 10.20.50.0/24 in the target data center. Hence when the virtual machine is migrated, its IP address is automatically changed to 10.20.50.40 on the target data center.

In the target data center, we have mapped the Production subnet to the rehearsal subnet. Hence when you perform the rehearsal operation, the rehearsal virtual machine is mapped on the rehearsal subnet.

For example, in your source data center you have two subnets (ProdSN1 10.20.30.0/24) with Production as the purpose and (ProdSN2 10.20.40.0/24) with rehearsal as the purpose. The Production subnet is mapped with the Rehearsal subnet. A similar setup is present on your target data center (RecovSN1 10.20.50.0/24 with Production as the purpose and (RecovSN2 10.20.60.0/24). Both the subnets having the purpose as Rehearsal are paired.

Note: Rehearsal subnet should be configured such that it is isolated from the production network.

You may want to perform Rehearsal operations in only one of the data centers. In that case, you can have network objects (for example, subnet) with purpose as Rehearsal in only that data center. If you map network objects with purpose as Rehearsal later, you need to edit the resiliency group with **Edit Configuration** or with **Customize Network** option.

Network pairs for recovering virtual machines to Google Cloud Platform (GCP)

You need to consider the below mentioned points before creating network pairs for recovering or migrating virtual machines to GCP :

- You can only configure a network interface when you create an instance.

- You cannot delete a network interface without deleting the instance.
- Each network interface configured in a single instance must be attached to a different VPC network, and each interface must belong to a subnet whose IP range does not overlap with the subnets of any other interfaces. The attached network can be a standalone VPC network or a Shared VPC network.
- For a workload multiple subnets having single VPC cannot be provisioned on GCP.
- For a workload multiple NICs in same subnet cannot be provisioned on GCP
- For shared subnets discovery, the host project name is appended to the VPC name to distinguish the shared subnets from the another project.

Figure 1-7 Overview of network structure for recovery to Google Cloud Platform

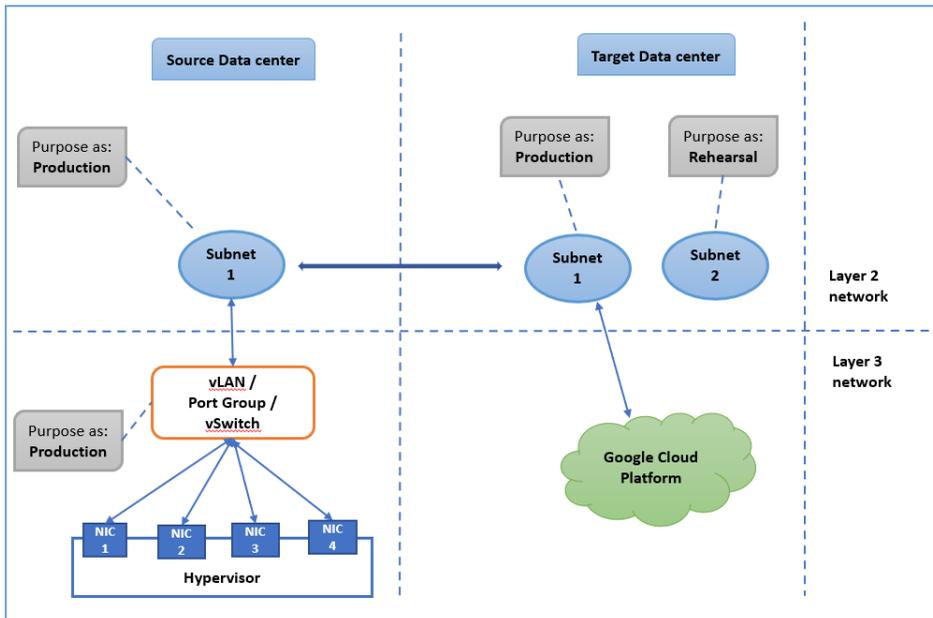


Table 1-65 Network Mapping types for Google Cloud Platform

Virtualization technology	Cloud environment	Network Mapping Type	Mandatory	Purpose
VMware	Google Cloud Platform	Subnet to Cloud Subnet	Yes	To connect the VNIC of the recovered virtual machine to the cloud network.
Hyper-V	Google Cloud Platform	Subnet to Cloud Subnet	Yes	To connect the VNIC of the recovered virtual machine to the cloud network.

[Creating network pairs between source and target data centers](#)

Creating network pairs between source and target data centers

The create network pair operation eliminates the need to manually connect each virtual machine to a network at the target data center. After you have paired the networks successfully, the target networks and the IP addresses are computed programmatically, and applied to the virtual machines.

Adding subnets using IPv6 address and pairing them across data centers is supported. You can pair subnets of same networks, that means pairing of subnets can be done for:

- IPv4 subnet to IPv4 subnet mapping
- IPv6 subnet to IPv6 subnet mapping

IPv6 network support is not applicable for cloud environments.

This step is a prerequisite for cloud recovery data center, if you want to override the default IP settings and customize the IP addresses when the virtual machines starts.

The mapping requirement depends on the target data center and is required to be done before you perform any disaster recovery operation. Before you create and map the network objects, ensure that all the assets are configured in the source as well as in target data center.

To create network pairs between production and recovery data centers

1 Navigate



Infrastructure Pairing (navigation pane)

2 Do one of the following:

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

3 In the **Network Mapping** page, select the data centers from source and the target data centers drop down to view the networks which are configured according to the vCenter or Hypervisor added .

Note that: Once you have configured hypervisor or cloud configuration then you will see the networks will be seen if they are discovered.

Networks Types options from the drop down will be displayed based on the virtualization technology or the cloud data centers configured.

4 If you want to view all the network objects, select the checkbox of **Show All Networks**. If you want to view only the network objects with purpose as **Production** , uncheck the checkbox.

5 Select the network objects from the lists and click **Move selected**.

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

Create resiliency groups

After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery. Refer following topics:

- Configuring a resiliency group for basic monitoring
- Prerequisites for configuring resiliency groups for recovery to Google Cloud Platform
- Configure resiliency groups for recovery to Google Cloud Platform

Configuring a resiliency group for basic monitoring

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

There are two types of pre-activated service objectives:

- Monitor assets - provides only monitoring, start, and stop operations
- Recover hosts - provides recovery operations as well as the start and stop operations

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

To manage assets for basic monitoring

1 Prerequisites

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

2 Navigate



Assets (navigation pane) **Unmanaged** tab > **Manage & Monitor Assets**

You can also launch the wizard from the **Overview** tab.

3 Select the assets:

- Select **Host** as the asset type, select the data center, type, and other filters as needed to display a list of assets.
- Drag and drop the selected assets to **Selected Instances**.

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

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

6 Supply a name for the resiliency group.

7 Verify that the new resiliency group is added to the **Resiliency Group(s)** tab.

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

Note: If the instances are created from BYOS image or there is licensing issues with instances, then Resiliency Platform operations may fail.

See [“About resiliency groups with assets”](#) on page 528.

Managing virtual machines for remote recovery (DR) to Google Cloud Platform

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

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

To manage virtual machines for remote recovery in Google Cloud Platform (GCP)

1 Prerequisites

See [“Prerequisites for configuring VMware virtual machines for recovery to Google Cloud Platform”](#) on page 206.

See [“Prerequisites for configuring Hyper-V virtual machines for recovery to Google Cloud Platform”](#) on page 518.

2 Navigate

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

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

3 Select the assets:

- Select the **Host** as Asset Type and select the other filters as needed to display a list of virtual machines.
- Drag and drop virtual machines or select the asset and click **Next**.

4 Next page displays the environment for the selected assets on **Review Environment** panel.

For resiliency group consisting of VMware virtual machines, ensure that each resiliency group is mapped to only one ESX cluster.

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

6 Select the target (recovery) data center on **Select Target Datacenter** page.

7 Review the information on the **Configure Resiliency Platform Data Mover** page and click **Next**.

- 8 Select the Replication Gateway pairs on **Select Replication Gateway pair** page.
- 9 Select the volume type.
See [“Volume type selection options”](#) on page 519.
- 10 In the **Confirm Resiliency Platform Data Mover Details** panel, verify the Replication Gateway pair selection and the asset information.
- 11 On the **Customize panel** perform the actions.
See [“Customize panel for Google Cloud Platform”](#) on page 520.
- 12 On the **Network Summary** panel, verify the asset name, MAC address and IP address of the assets in the resiliency groups and click **Next**.
- 13 Complete the network customization steps for the virtualization technology on the **Customize Network** panel.
See [“Network customization options”](#) on page 523.
- 14 Use the **Customize System Generated workflows** panel to enable manual intervention at predefined points during the Migrate and Recover operations on both the data centers. This is an optional step. See [“About manual intervention”](#) on page 525.
- 15 Verify the summarized information and enter a name for the resiliency group and click **Submit**.

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

Prerequisites for configuring VMware virtual machines for recovery to Google Cloud Platform

Before you run the wizard to configure disaster recovery protection for a resiliency group of VMware virtual machines, ensure that you have met the following prerequisites for the virtualization environment:

- VMware Tools must be installed on the virtual machines.
- If you add new disks, ensure that they are visible from the guest operating system.
- Ensure that all the ESXi servers of the cluster belong to the same VMware data center.

- If a virtual machine has more than one ethernet adapter, then all of them should have either static IP configuration or DHCP IP configuration. A mix of static and DHCP IP configuration is not supported on the same virtual machine.
- All the virtual disks must be connected to the virtual SCSI controllers. Other controller types are not supported.
- The datastores on which the virtual machine disks reside must be accessible to the Replication Gateway on the production data center.
- Ensure that the configuration files (vmx files), residing on the datastores, are not located inside a folder. Also the configuration files must not have any special characters in their names.
- Though special characters are allowed in a password, you cannot use space in the password.
- Configuring resiliency group for remote recovery using Data Mover fails if a virtual machine has snapshots. Hence remove all snapshots before proceeding with the operation.
- Enable the UUID for the virtual machines (disk.enableuuid=true). Refer to VMware documentation for details.
- Ensure that the Replication Gateways have sufficient storage to handle the replication for the planned number of protected virtual machines. Both the on-premises gateway and the cloud gateway must have external storage equivalent to 6GB for each asset protected by the gateway pair. For example, if a gateway pair supports 10 virtual machines, the on-premises gateway and the cloud gateway must each have 60 GB of external storage.
- On Linux virtual machines, ensure that the NIC name in the file matches with the actual NIC name on the system.
- On Linux virtual machines, ensure that all the virtual NICs attached to the selected virtual machines have a valid IP address. The IP address may be assigned statically or via DHCP protocol.
- If the status of the virtual machine on the recovery data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.

Note: Resiliency Platform does not support protecting a virtual machine which is created by cloning another VMware virtual machine that is already protected under Veritas Resiliency Platform.

See [“Managing virtual machines for remote recovery \(DR\) to Google Cloud Platform”](#) on page 517.

Volume type selection options

This panel is displayed when you are configuring your virtual machines for recovery to AWS or Google Cloud Platform.

Table 1-66 Volume type options available for AWS

Options	Description
Volume Type	Select the volume type for the disks. You can apply the selections either to all the virtual machines or customize for each.
IOPS (This field is enabled if Provisioned IOPS SSD option is selected)	Enter the IOPS required if the volume type is Provisioned IOPS SSD. Note: Refer to AWS documentation for more information on IOPS permitted for specific volume type and size.
KMS Encryption Key	Select the KMS Encryption Key, if you want to create encrypted volume in AWS, before selecting KMS key. Ensure that KMS key has all the required permissions. Note: Refer to AWS documentation for KMS key permissions and IAM role attached to IMS.
Availability Zone	The availability zones which are listed are based on the Replication Gateway pairs that you have chosen. You can apply the selected volume type and available zone either to all the virtual machines or customize for each.

Table 1-67 Volume type options available for Google Cloud Platform

Options	Description
Volume Type	Select the volume type for the disks. You can apply the selections either to all the virtual machines or customize for each. Note: The Regional Volume Type is not applicable for Boot disk. Even though the Regional Volume Type is selected in Apply All action, this volume type is not applicable for Boot disk.
IOPS (This field is enabled if Extreme Persistent Disk option is selected)	Enter the IOPS required if the volume type is Extreme Persistent Disk. Note: Refer to GCP documentation for more information on IOPS permitted for specific volume type and size.

Table 1-67 Volume type options available for Google Cloud Platform
(continued)

Options	Description
Encryption Key	<p>Select the Encryption Key, if you want to create encrypted volume in GCP. Before selecting key, ensure that key has all the required permissions.</p> <p>Note: Refer to Google Cloud Platform documentation troubleshooting section for the actual command in Resiliency Platform Product documentation.</p> <p>All the disks in Google Cloud Platform are already encrypted by Google Managed Key. Hence, you can choose to create and manage your own keys. These keys are also applicable for all or selected disks of a workload.</p>
Zone	The zones which are listed are based on the Replication Gateway pairs that you have chosen. You can apply the selected volume type and available zone either to all the virtual machines or customize for each.
Replica Zone	This field is enabled when Regional disks types are selected.

Customize panel for Google Cloud Platform

This panel is displayed when you are configuring your virtual machines for recovery to Google Cloud Platform.

Table 1-68 Select Attributes panel

Options	Description
Use same firewall for rehearsal operation similar to the one in production checkbox	<p>Checkbox is selected by default.</p> <p>Machine type and Target VM values are already present.</p> <p>If the checkbox is uncheck, provide the following details:</p> <ul style="list-style-type: none"> ■ Network tag or Rehearsal Network Tag
Target VM Name	Value is already present. It can be changed later.
Machine Type dropdown	Select the value.
Network tag or Rehearsal Network Tag	These values are already present when the Use same firewall for rehearsal operation similar to the one in production checkbox is selected.

Select NIC panel

While configuring your virtual machines in Goggle Cloud Platform, each network interface of an instance must be attached to different VPC and also each network interface must belong to a different subnet. If network mapping conflicts to above

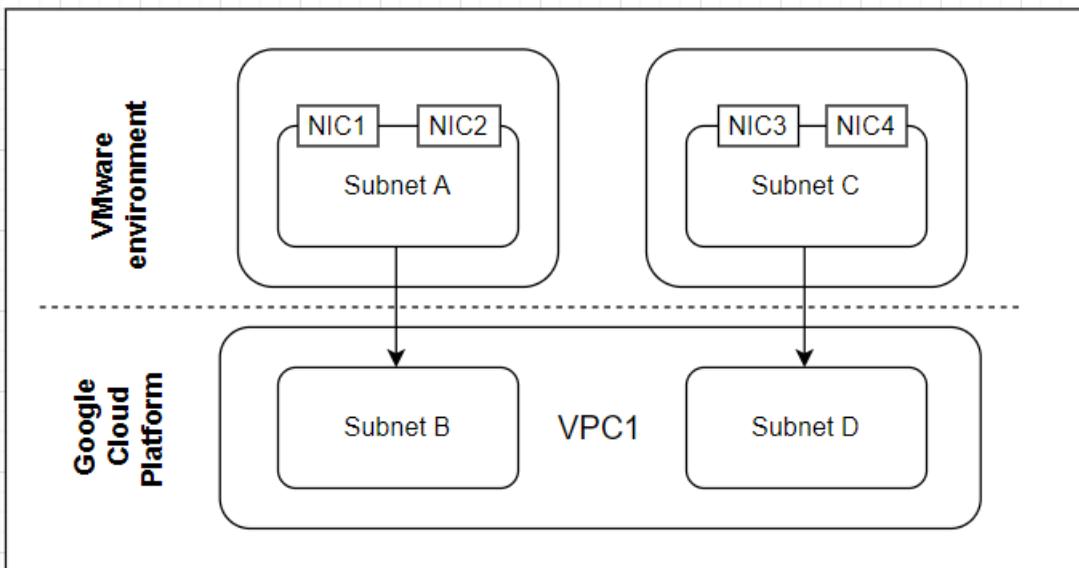
requirement, then you must select **subnet** from dropdown and within the subnet select the primary NIC from the dropdown on **Select Primary NIC** panel and click **Next**.

Below is the example of network mapping which conflicts with Google Cloud Platform requirement:

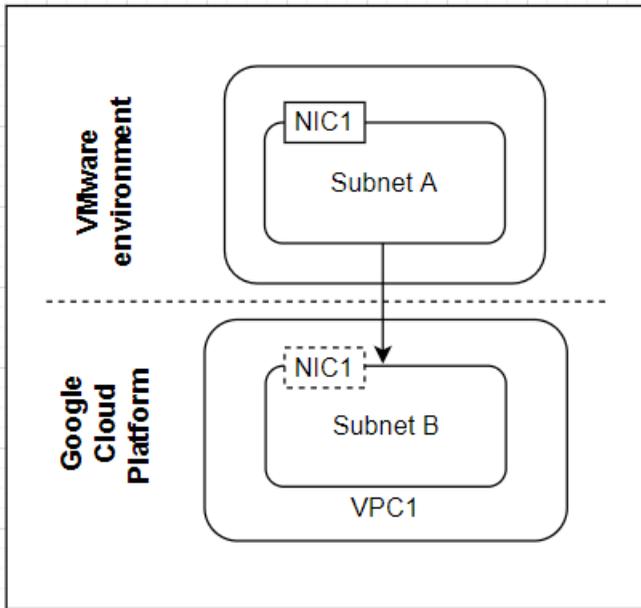
If you create network mapping from VMware or Hyper-V to Google Cloud Platform, say subnet A to B and subnet C to D and there are two NICs which exist in subnet A and two NICs exists in subnet C. On Google Cloud Platform, subnet B and subnet D belongs to same VPC say VPC1.

Below is the diagrammatical representation of the configuration explained:

Figure 1-8 Diagrammatical representation of the network configuration:



So considering the above configuration, the mapping in Google Cloud Platform should have each network interface attached to different VPC and so for different subnet there must be one NIC created out of above four NICs. Hence, as described in the below diagram if you select a subnet say Subnet A and then select a primary NIC say NIC1 from subnet A using dropdown in **Select Primary NIC** panel, after migrating the virtual machines the NIC1 is created on Subnet B in GCP environment.

Figure 1-9 Example of the configuration

Network customization options

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

To customize the static IP of Windows guest virtual machines in the VMware environment following are the two options:

- Use global user credentials for IP customization of Windows virtual machines. This option uses the Windows global user credentials. These credentials must be configured in advance.
- Install IP customization service on Windows virtual machines. After you finish installing the IP customization service, ensure that the following settings are disabled:
 - User Account Control: Admin Approval mode for the Build-in Administrator account
 - User Account Control: Run all administrators in Admin Approval mode. This option installs a service in the virtual machine to assist in IP customization. This service does not have any other functionality and does not require any

inbound or any outbound communication. The Resiliency Platform does not store the credentials which are provided to install the IP customization service. This option does not enforce authentication domain reachability on target data center unlike the global user option's requirement. This option is only applicable for VMware environment with third-party replication and Resiliency Platform Data Mover with VAIO framework.

You can do the following in this panel:

- Choose between the following two options for IP customization of Windows Virtual Machines
 - Use global user credentials for IP customization of Windows virtual machines.
 - Install IP customization service on Windows virtual machines.
- Manage PTR records
 - For Windows DNS and for Linux Bind, if you want Resiliency Platform to customize DNS settings, then DNS records should not exist at the target data center.
 - For Windows DNS, if you configure a user in Resiliency Platform for DNS customization, then that user should also have rights to update DNS records added by any other user on the DNS.
- Choose to continue with DR operations even if DNS updates fail.

You can customize the IPs for Production and Rehearsal networks. Customizing the IPs of a virtual machine overrides the default IP settings when the virtual machine starts at the target data center.

If the subnet mapping is already done, and the mapped subnets are of equal mask, then the computation of projected static IP is done based on the subnet mappings by Resiliency Platform. This projected IP address for the target data center can be edited. If the mapped subnets are of unequal masks, then you need to enter the IP address manually.

If the subnet mapping is not done, then you can either select a subnet from the drop-down list and apply it to all the target IP addresses or you can select a separate subnet for each target IP address. You also need to enter the IP address for the target data center. Since IPv6 network support is enabled and subnets can be created using IPv6 address, you can see the IPv6 subnets in the drop-down. Either you can apply the IPv6 address to all the target IP address or you can select a separate subnet for each target IP address.

You can choose to continue with the DR operation if the IP customization fails. Note that this is possible only if the virtual machines have static IPs.

To customize the static IP of Windows guest virtual machines in the VMware environment, Resiliency Platform requires the user name and password to log on to the Windows virtual machines. To configure this user name and password go to **Product Settings > User Management > Windows Global User**.

See “[Configuring Windows global user](#)” on page 482.

If the target data center is in cloud, then ensure that the IPs used for network customization are not already in use on the cloud.

If you choose to apply DNS customization, then you can add a host name to IP mapping of the DNS.

About manual intervention

Resiliency Platform lets you pause the Migrate and Recover operation at certain predefined stages. This gives you time to perform any manual tasks on the assets if required. On completion of your tasks, you can resume the operation from **Activities** or **Recent Activities** menu.

When you are configuring a resiliency group for disaster recovery (DR), one of the panels lets you select the pause or the manual intervention point. Similarly when you are configuring a Virtual Business Service (VBS) for DR you can choose the manual intervention points.

Following are the predefined points while configuring a resiliency group:

- After stopping the assets on the source data center.
- Before registering the assets on the target (recovery) data center.

While configuring a VBS, the predefined points are during stop of a tier and start of a tier.

- Stop of a tier: Before stopping and after starting the workloads.
- Start of a tier: Before stopping and after starting the workloads.

Enabling manual intervention is an optional step. This panel is displayed when you are configuring the resiliency group for the following scenarios:

- Recovery to on-premises data center using Resiliency Platform Data Mover.
- Recovery to on-premises data center using any 3rd party replication technology.
- Recovery to any cloud data center.

Note: The Migrate and Recover operations remain in paused state till you resume the operations.

You can view the time spent for manual intervention in the following reports:

- Activity Distribution History
- Recovery Activity History by RG
- Recovery Activity History by VBS

See [“Managing a running activity”](#) on page 582.

Advanced features

Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.

- See [“About virtual business services”](#) on page 526. Managing virtual business services
- See [“About resiliency groups with assets”](#) on page 528. Managing resiliency plans
- See [“About evacuation plan”](#) on page 530. About evacuation plan

About virtual business services

For a business service to work properly, it is important that all of its tiers and components are up and working together. From a business continuity point of view, it is important to not just ensure that individual tiers are up and running but also the entire business service.

A virtual business service (VBS) is a logical collection of resiliency groups that function together to perform a particular business service. A VBS enables easy management of multi-tier business services. For example, you can group a web server resiliency group, a database resiliency group, and a payroll business logic resiliency group into a VBS called `payroll`. You can start, stop, monitor, manage, or recover that VBS as a single entity.

Note: If a VBS consists of resiliency groups that are in maintenance mode, then you cannot perform any operations on the VBS.

Understanding tiers

Within a VBS, resiliency groups are arranged in tiers. Tiers represent the logical dependencies between the resiliency groups and determine the relative order in which the resiliency groups start and stop. For example, the database resiliency group must start before the application server resiliency group and the web server resiliency group, so the database resiliency group must go in the lowest tier. The

application server resiliency group must start after the database resiliency group, so it goes in the next tier. The web server resiliency group must start last, so it goes into the top tier. Later, if you add a resiliency group to the VBS, you can manage it as part of the IT service by placing it in the appropriate tier.

Various configuration states of virtual business services

It may happen that during upgrade or configuring VBS, the workflow may stuck at some particular state where manual reconfiguration is required.

Table 1-69 VBS configuration states and resolution

Configuration states	Virtual business service may lead into the state on following events	Resolution
Configuring	<ul style="list-style-type: none"> ■ Create virtual business service ■ Edit virtual business service ■ Edit resiliency group associated with virtual business service ■ Delete resiliency group associated with virtual business service 	N/A
Configuration failed	<ul style="list-style-type: none"> ■ Create virtual business service fails ■ Edit virtual business service fails ■ Virtual business service reconfiguration triggered by following event fails: <ul style="list-style-type: none"> ■ Edit resiliency group associated with virtual business service ■ Delete resiliency group associated with virtual business service 	Edit the virtual business service manually.

Table 1-69 VBS configuration states and resolution (*continued*)

Configuration states	Virtual business service may lead into the state on following events	Resolution
Unconfiguring	Delete virtual business service	N/A
Unconfiguration failed	Delete virtual business service fails	Retry deleting the virtual business service.
Refreshing	Refresh virtual business service	N/A
Reconfiguration pending	<p>Virtual Business Service is currently in Configuring/Refreshing state and at the same time Reconfiguration of virtual business service is triggered by event:</p> <ul style="list-style-type: none"> ■ Edit resiliency group associated with virtual business service ■ Delete resiliency group associated with virtual business service 	<p>N/A</p> <p>Note: This is an intermediate state and should get resolved automatically.</p>

See [“About virtual business services”](#) on page 526.

About resiliency groups with assets

Resiliency groups are the unit of management and control in Veritas Resiliency Platform. After assets are added to Resiliency Platform, you organize related assets into a resiliency group that you can protect and manage as a single entity. A resiliency group can have only physical machines or only virtual machines, a mix of physical and virtual machines is not supported. Similarly it can contain either all applications or all InfoScale applications but not a mix of both.

For example, you can organize several physical or virtual machines into a resiliency group, and name it `VM_Finance`. When you perform an operation on the `VM_Finance` resiliency group using the Resiliency Platform console, the operation is performed on all the assets that belong to the resiliency group. For example if you run the Start operation on the resiliency group, all the assets (physical machines or virtual

machines) that belong to the resiliency group start booting. Or if you perform any of the disaster recovery operations such as Migrate on the resiliency group, all the assets within the group are migrated to the selected target data center.

Ensure that the following prerequisites are met while creating a resiliency group for remote recovery:

- Ensure that all the virtual machines that are to be grouped in a single resiliency group belong to a single hypervisor or virtualization server (if not clustered), or to a single cluster.
- Encryption and compression are disabled on Hyper-V servers.
- Refresh the virtualization server (vCenter server or Hyper-V server), the host, and the cloud discovery before proceeding to create the resiliency group.
- If the recovery is on Azure, then ensure that the virtual machine names should always start and end with alphanumeric characters. The name can contain periods (.), hyphens(-), or underscores(_) in the middle.
- If you are using third-party replication, ensure that the assets consume storage from the same consistency groups. E.g. EMC SRDF device group, NetApp Volume, 3PAR replication group, and so on.

The operations available for a resiliency group depend on how it is configured. While configuring a resiliency group, you need to select a service objective. If you select a service object that supports remote recovery, then you can perform disaster recovery operations such as Migrate and Take over on the resiliency group.

Optionally you can use a service objective that only monitors the assets or the applications and provides only basic operation capabilities like start and stop operations and no remote recovery operations. Using the Copy service objective, you can recover the virtual machines from NetBackup generated backup images to the target data center.

See [“Configuring a resiliency group for basic monitoring”](#) on page 529.

Configuring a resiliency group for basic monitoring

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

There are two types of pre-activated service objectives:

- Monitor assets - provides only monitoring, start, and stop operations
- Recover hosts - provides recovery operations as well as the start and stop operations

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

To manage assets for basic monitoring

1 Prerequisites

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

2 Navigate



Assets (navigation pane) **Unmanaged** tab > **Manage & Monitor Assets**

You can also launch the wizard from the **Overview** tab.

3 Select the assets:

- Select **Host** as the asset type, select the data center, type, and other filters as needed to display a list of assets.
- Drag and drop the selected assets to **Selected Instances**.

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

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

6 Supply a name for the resiliency group.

7 Verify that the new resiliency group is added to the **Resiliency Group(s)** tab.

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

Note: If the instances are created from BYOS image or there is licensing issues with instances, then Resiliency Platform operations may fail.

See [“About resiliency groups with assets”](#) on page 528.

About evacuation plan

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

Using the evacuation plan template you can define the sequence in which the virtual business services (VBS) should be migrated from the production data center to the recovery data center. Resiliency groups that do not belong to any VBSs, are appended at the end of the evacuation plan workflow after the VBS. If there are large number of VBSs then up to 5 VBSs within a priority group are migrated in

parallel to the recovery data center. Similarly, if there are large number of resiliency groups, up to 10 resiliency groups are migrated in parallel.

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

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

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

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

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

For a VBS or a resiliency group to successfully evacuate to the target data center, it should meet the following criteria:

- VBS or resiliency group that belong to the evacuation plan must be configured for disaster recovery.
- VBS can contain resiliency groups some of which are configured for disaster recovery and some using the service objective with data availability as Copy.
- Resiliency group must belong to only one VBS.

When you generate a plan, an appropriate warning is shown listing the assets that are excluded from the plan.

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

- Evacuate
- Rehearse evacuation
- Cleanup evacuation rehearsal
- Regenerate

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

- VBSs are added, modified, or deleted.
- Resiliency groups are added and configured for disaster recovery.
- Resiliency groups which were configured for disaster recovery are deleted.

- Existing resiliency group is configured for disaster recovery.

No action is required in the following scenarios:

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

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

Perform remote recovery operations

Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.

- See [“Performing the rehearsal operation for virtual machines”](#) on page 532. Performing the rehearsal operation for virtual machines
- See [“Performing cleanup rehearsal for virtual machines”](#) on page 535. Performing cleanup rehearsal for virtual machines
- See [“Migrating a resiliency group”](#) on page 536. Migrating a resiliency group
- See [“Performing the resync operation for virtual machines”](#) on page 537. Recovering resiliency group of virtual machines
- See [“Recovering a resiliency group using replication-based recovery”](#) on page 539. Performing the resync operation for virtual machines

Performing the rehearsal operation for virtual machines

Use the **Rehearsal** option on the Resiliency Platform console to ensure the disaster recovery readiness of the assets in your protected resiliency groups.

From version 3.5, Rehearsal operation can be performed on Azure or (Azure Stack) data center when it is deployed as source as well as a target data center.

For recovery on AWS cloud:

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

during Rehearsal is significantly reduced. Which reduces the overall time taken to complete the operation.

Note: This setting is specific to recovery of virtual machines from VMware to VMware if Resiliency Platform Data Mover is used and recovery of physical machines to VMware data center:

While performing the rehearsal operation, DRS automation level for the target Replication Gateway should be set to manual or it should be disabled.

To perform the rehearsal operation

1 Prerequisites

See [“Prerequisites for rehearsal operation for virtual machines”](#) on page 534.

2 Navigate



Assets (navigation pane) > **Resiliency Group(s)** tab

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

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

Rehearsal operation on virtual machines for CDP

The rehearsal option on the Resiliency Platform console is used to ensure the disaster recovery readiness of the assets in your protected resiliency groups. You can perform the rehearsal operation for the resiliency groups where CDP is enabled.

To perform rehearsal operation for virtual machines for CDP

Navigate

1



Assets (navigation pane) > **Resiliency Group(s)** tab

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

3 Select the target data center and then select the recovery points to perform the rehearsal operation.

4 Click **Next**.

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

See [“Performing cleanup rehearsal for virtual machines”](#) on page 535.

Prerequisites for rehearsal operation for virtual machines

Before you run the rehearsal operation for a resiliency group, ensure that you have met the following prerequisites:

- For VMware virtual machines, ensure that the datastores have enough free space for the swap files for the on-premises virtual machines and the virtual machines created by the rehearsal operation on the recovery data center. The size of the swap files is same as that of the virtual machine memory size.
- For VMware virtual machines, ensure that the mapping of all the required port groups across the data centers is complete.
For Hyper-V virtual machines, ensure that the mapping of all the required virtual switches across the data centers is complete.
See [“Creating network pairs between source and target data centers”](#) on page 514.
- Each type of replication has prerequisites and limitations for the rehearsal operation.
- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs. This is required if the recovery data center is AWS, Azure cloud.
- If the recovery data center is in AWS, then configure a rehearsal subnet in the cloud. The rehearsal and production subnet should be in the same VPC.
- If the recovery data center is in cloud, then you need to set the SAN policy to either OnlineAll or OfflineShared based on whether you have shared or non-shared disks. For more details refer to [Microsoft Documentation](#).
- If the status of the virtual machine on the recovery data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.
- When the **Use Same WWN** option is checked, Resiliency Platform needs to do additional storage operations to map or unmap the LUNs, swap WWNS of LUNs during the workflow and that requires extra rescans and validations. These operations are not required when the option **Use Same WWN** is disabled. This option can be set only for when the DR operation is invoked from Resiliency Platform. It can be reset once the DR activity is complete.

Note: This setting is specific to recovery of virtual machines from VMware to VMware if Resiliency Platform DataMover is used and recovery of physical machines to VMware data center:

While performing the rehearsal operation, DRS automation level for the target Replication Gateway should be set to manual or it should be disabled.

See [“Performing the rehearsal operation for virtual machines”](#) on page 532.

Performing cleanup rehearsal for virtual machines

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

Note: Any snapshots of the cloud volumes that are taken external to Veritas Resiliency Platform may cause failure in rehearsal cleanup. This is applicable only for Orange Recovery Engine.

This setting is specific to recovery of virtual machines from VMware to VMware if Resiliency Platform data Mover is used and recovery of physical machines to VMware data center: While performing the rehearsal operation, DRS automation level for the target Replication Gateway should be set to manual or it should be disabled.

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

- A separate copy of virtual machine when you use Hyper-V Replica for data replication.
- A new registered virtual machine that has its virtual machine data files (VHDX) residing on snapshot LUNs when array-based replication (for example, EMC SRDF) is used for data replication.

Using NetBackup

When your assets are configured for remote recovery using a service objective where the data availability mode is Copy, then during the rehearsal operation virtual machines are created on the recovery data center with the selected backup image. These virtual machines and the data are deleted during the cleanup operation.

To perform cleanup rehearsal

- 1 Navigate to **Assets** (navigation pane) > **Resiliency Group(s)** tab.
- 2 Double-click the resiliency group to view the details page. Click **Cleanup Rehearsal**.
- 3 Select the target data center, and click **Next**.

See [“Performing the rehearsal operation for virtual machines”](#) on page 532.

Migrating a resiliency group

Migration refers to a planned activity involving graceful shutdown of physical and virtual machines at the source data center and starting them at the target data center. In this process, replication ensures that consistent data of the assets is made available at the target data center which could be the on-premises or cloud data center. In Veritas Resiliency Platform, the migration of assets is achieved by grouping them in a resiliency group, configuring disaster recovery for the resiliency group, and thereafter performing the migrate operation on this resiliency group.

Consider the following:

- If you perform the recover operation, then you must perform the Resync operation before you migrate back to the production data center.
- If the **Enable reverse replication** option is not selected, then before migrating the virtual machines to the target data center and after migrating back to the source data center, you need to perform the Resync operation. See [“Performing the resync operation for virtual machines”](#) on page 537.
- If the recovery data center is Azure cloud, then after you migrate from the cloud data center to the on-premises data center, you need to refresh Azure cloud to rediscover the cloud-based objects.
- During the migrate operation, virtual machines on the source data center are gracefully shut down. If you manually shut down the virtual machine before performing the migrate operation, and if the shut down was not graceful, then the migrate operation may fail. This is applicable when the replication technology is Resiliency Platform Data Mover and the target data center is in cloud.
- When you upgrade from an earlier version to version 3.2 or later, then after performing the migrate operation, a risk is raised. This risk is regarding the changes in NIC configuration when you migrate to any cloud data center. Suppress this risk while the resiliency group is online on cloud. Migrate back to the on-premises data center and then edit the resiliency group to fix the NIC configuration.

Prerequisites

- Ensure that the Data Mover connection status is **Connected**, Data State is **Consistent**, and Replication State is **Active**.
- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs. This is required if the recovery data center is AWS, Azure or Google Cloud Platform.
- For VMware virtual machines, ensure that the network mapping of all the required port groups, or subnets across the data centers is complete.

For Hyper-V virtual machines, ensure that the network mapping of all the required virtual switches across the data centers is complete.

See “[Creating network pairs between source and target data centers](#)” on page 514.

- If the recovery data center is in AWS, Azure or Google Cloud Platform then ensure that the network mapping of all the required subnets across the data centers is complete.
- If the recovery data center is in cloud, then you need to set the SAN policy to either OnlineAll or OfflineShared based on whether you have shared or non-shared disks. For more details refer to [Microsoft Documentation](#).
- If the status of the virtual machine on the recovery data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.
- For the replication technology HPE 3PAR Remote Copy, ensure that for VMware virtual machines the `config.vpxd.filter.hostRescanFilter` value is set to false.

To migrate a resiliency group

1 Navigate



Assets (navigation pane) > **Resiliency Group(s)** tab

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

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

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

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

For more information on troubleshooting specific scenarios,

Performing the resync operation for virtual machines

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

Since the production data center is not working, the data replication between the two sites does not happen. After the production site is back up and running, you

need to prepare the production site for the next failover or for a migration operation. This preparation includes cleaning up any residue and resuming the replication from the recovery to the production site.

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

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

Consider the following if you have configured your assets for recovery to vCloud Director without adding the VMware vCenter server or Hyper-V server:

You need to perform the resync operation after Migrate or Recover. In the **Activities** panel, if the workflow is in **Paused** state for "Resync Replication" subtask, then you need to manually start the physical and virtual machines. Ensure that the physical and virtual machines boot from PXE OS of the replication gateway that is configured as PXE server. You can verify the virtual machines boot progress from vCenter or Hyper-V console by checking the virtual machine console. When the boot is complete, click **Resume** so that the workflow proceeds with synchronizing data from the disks on recovery datacenter to those on the on-premises data center. After Resync operation is complete, do not shut down the virtual machines, otherwise the subsequent Migrate or Recover operations fail.

Performing the resync operation

1 Prerequisites

- If the target (recovery) data center is on-premises, and the last performed operation was recover, then you may need to restart the Hyper-V server or the ESX server. Although the Resync operation cleans up any residue on the source data center before resuming replication, there could be some residues that can be cleaned up only by restarting the hypervisors.
- For the replication technology HPE 3PAR Remote Copy, ensure that for VMware virtual machines the `config.vpxd.filter.hostRescanFilter` value is set to false.
- If the status of the virtual machine on the recovery data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.

Prerequisites for resync operation of physical workloads

- The resiliency group must be configured for DR and migrated to the recovery datacenter.

- Ensure that PXE boot server is configured on the on-premises Replication Gateway.
 - Ensure to set network boot (PXE boot) as the first boot priority in the system BIOS.
 - If you are not using 3rd party DHCP server, then configure DHCP server on the PXE Boot server that is configured on the on-premises Replication Gateway.
- 2 Navigate to **Assets** (navigation pane) > **Resiliency Group(s)** tab
 - 3 Double-click the resiliency group to view the details page. Click **Resync**.
 - 4 In the **Resync** panel, select the production data center name from the drop-down list, and click **Next**.

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

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

After completing recover from cloud and resync, the resiliency group details page shows entries for deleted or unavailable virtual machines on cloud data center. To remove these stale entries, after resync is complete, edit the resiliency group with **Edit Configuration** intent. You may submit the wizard without making any changes.

Recovering a resiliency group using replication-based recovery

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

Perform the resync operation after successful completion of recover operation.

If you are recovering to vCloud Director data center, without adding Hyper-V Server or vCenter Server, then recover operation from cloud to production (on-premises) data center is not supported.

Prerequisites

- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs. This is required if the recovery data center is AWS, Azure cloud.
- For VMware virtual machines, ensure that the network mapping of all the required port groups, or subnets across the data centers is complete.
For Hyper-V virtual machines, ensure that the network mapping of all the required virtual switches across the data centers is complete.
See [“Creating network pairs between source and target data centers”](#) on page 514.
- If the source data center is in AWS, then ensure that the network mapping of all the required subnets between the source and target data center is complete.
- There should not be any resources on Azure having the same name or substring of name as that of the virtual machine display name or FQHN name on on-premises data center.
- If the source data center is in cloud, then you need to set the SAN policy to either OnlineAll or OfflineShared based on whether you have shared or non-shared disks. For more details refer to [Microsoft Documentation](#).
- If the status of the virtual machine on the source data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.
- For the replication technology HPE 3PAR Remote Copy, ensure that for VMware virtual machines the `config.vpxd.filter.hostRescanFilter` value is set to false.

When you upgrade from an earlier version to version 10.0 or later, then after performing the recover operation with replicated data, a risk is raised. This risk is regarding the changes in NIC configuration when you recover to any cloud data center. Suppress this risk while the resiliency group is online on cloud. Migrate back to the on-premises data center and then edit the resiliency group to fix the NIC configuration.

To perform recover operation on virtual machines

1 Navigate



Assets (navigation pane) > **Resiliency Group(s)** tab

- 2 Double-click the resiliency group to view the details page. Click **Recover**.
- 3 Do the following:

- Select the target data center.
- If there is an outage on the source data center, select the **Confirm outage of assets** check box.
- During the recover operation, if Resiliency Platform detects a probability of data loss, you have the option to abort the recover operation to avoid any data loss. Select the check box if you want to abort the operation in such a situation.
- Click **Continue** for warnings.

4 Click **Submit**.

To perform recover operation on resiliency group configured with CDP enabled

1 Navigate



Assets (navigation pane) > **Resiliency Group(s)** tab

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

3 To select the recovery point for CDP, do the following:

4 On the **Select Recovery points** panel, you can select the date, time, and range to list the recovery points with the latest data against the assets.

- Select the target data center.
- If there is an outage on the source data center, select the **Confirm outage of assets** check box.
- Do not select the **Abort recover if these subtasks fail** check box and click **Next**.
- Select the recovery points. To choose from a specific time range, select the date and the time range. Then select the hours or minutes since the start time. Click **Search**.

5 Click **Submit**.

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

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

For more information on troubleshooting specific scenarios,

Monitor assets

You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.

- See [“About risks”](#) on page 542.About risks
- See [“About reports”](#) on page 576.About reports
- See [“Managing a running activity”](#) on page 582.Managing activities

About risks

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

Risk Insight detects the changes to the state and configuration of your protected assets. It identifies if there is a risk to the recoverability or continuity of your protected assets. A periodic or schedule scan will generate the risk on the Resiliency Platform components. A periodic scan of every 30 minutes is run by the scheduler.

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

Risk insight generates two types of reports:

- **Current risk reports:** Provides the summary and detail information about all the current risks in your data center.
- **Historical risk reports:** Provides a summary and a detailed analysis of information about the risks in your environment during the specified period.

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

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

Table 1-70 Risk types

Risk category	Description
Recoverability	Risks that may impact the ability to recover and run the application on the recovery site.

Table 1-70 Risk types (*continued*)

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

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

Table 1-71 Risk types

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

From 3.2 onwards, you can probe and suppress a risk based on which component of Resiliency Platform the risk has occurred.

Probing a risk is way of evaluating a risk to check whether the risk is eliminated or still exists at the Resiliency Platform component. You cannot probe all the risks in the risk view. The risks which cannot be probed have details about the risk resolution. You can perform the given steps and probe the risk again. You have to wait for some time the to see the apply the changes.

For example, an IMS is disconnected due to network issues. The risk is then raised by the Resiliency Platform. When you probe this risk, details of this risk display what steps can be done to resolve this risk. There are some risks which cannot be probed. For those risks, you have to fix the risk and the probe it accordingly.

When you probe a risk, the details panel also displays a link to the proposed resolution for the risk. On clicking the link, the risk pop up is closed and you are redirected to the respective resolution page.

You can avoid a risk using suppress option on the risk view. You can suppress a risk for specific time; for few minutes or hours but cannot suppress a risk for indefinite period. If a risk is active after the suppress period is over, then it will appear under the Active Risks view. You can probe that risk or you can again suppress it for some time. Suppress option is available for all the risk in the Resiliency Platform. Suppressed risks will be seen in Risk view > Suppressed Risks.

Note: Risks are not generated for resiliency groups that are in maintenance mode. If a risk was raised before the resiliency group was placed in maintenance mode and that risk persists after exiting the mode, then the risk is shown.

See [“Predefined risks in Resiliency Platform”](#) on page 544.

Predefined risks in Resiliency Platform

[Table 1-72](#) lists the predefined risks available in Resiliency Platform. These risks are reflected in the current risk report and the historical risk report.

Table 1-72 Predefined risks

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
vCenter Password Incorrect	Checks if vCenter password is incorrect	15 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop operations ■ On secondary site: migrate or recover operations 	In case of a password change, resolve the password issue and refresh the vCenter configuration
VM tools not installed	Checks if VM Tools are not Installed. It may affect IP Customization and VM Shutdown	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Stop 	<ul style="list-style-type: none"> ■ In case of VMWare, install VMWare Tools ■ In case of Hyper-V, install Hyper-V Integration Tools
Snapshot reverted on Virtual Machine	Checks if snapshot has been reverted on virtual machine	5 minutes	Error	Resiliency Platform Data Mover replication	Perform the Resync operation on the resiliency group.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Resiliency Platform Data Mover daemon crashed	Resiliency Platform Data Mover filter is not able to connect to its counterpart in ESX. The replication process has stopped and is at risk	5 minutes	Error	Resiliency Platform Data Mover replication	<ul style="list-style-type: none"> ■ To continue the replication, you can move (VMotion) the virtual machine to a different ESX node in the cluster. ■ Troubleshoot the issue with this ESX node or raise a support case with Veritas.
DataMover virtual machine in no-op mode	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Veritas Replication policy has been detached	Veritas Replication policy has been detached from the disk associated with virtual machine.	5 minutes	Error	Migrate	Perform Resync operation on the affected resiliency group.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Asset disk configuration changed	Checks if disk configuration of any of the assets in the resiliency group has changed.	30 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal 	<p>Refresh the respective hosts, vCenter servers or Hyper-V servers and the cloud discovery. After refresh, probe the risk.</p> <p>After performing the above mentioned step even if the risk still exists, edit the resiliency group to first remove the impacted virtual machine from the resiliency group and then add it back to the resiliency group.</p>
Asset NIC configuration changed	Checks if NIC configuration of any of the assets in the resiliency group has changed.	30 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Resync 	<p>If the resilience group is online on the target data center, then either revert the NIC changes done on the virtual machines or suppress the risk to be able to migrate the assets back to the source data center. If the resiliency group is online on source data center, edit the resiliency group with Edit Configuration or Customize Network option to update the NIC configuration.</p>

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Invalid NIC Configuration	One or more NICs on the host are not configured properly.	Real time, while creating resiliency group	Error	Create resiliency group	Ensure that the keys NAME, DEVICE and HWADDR have appropriate values as per the details of each NIC in its configuration file.
Global user deleted	Checks if there are no global users. In this case, the user will not be able to customize the IP for Windows machines in VMware environment	Real time	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Edit the resiliency group or add a Global user
Failure to validate Windows Global User credentials for IP customization	This risk is raised if: <ul style="list-style-type: none"> ■ Windows Global User is not configured. ■ Windows Global User does not have appropriate credentials. ■ Virtual machine is offline while configuring resiliency group for disaster recovery. 	After the resiliency group is configured for disaster recovery	Warning	<ul style="list-style-type: none"> ■ Rehearsal ■ Migrate ■ Recover 	Add Windows Global Users with appropriate credentials. Edit the resiliency group using the Network Customization option to resolve the risk.
Missing heartbeat from Resiliency Manager	Checks for heartbeat failure from a Resiliency Manager	5 minutes	Error	All	Fix the Resiliency Manager connectivity issue
Infrastructure Management Server disconnected	Check for Infrastructure Management Server(IMS) to Resiliency Manager(RM) connection state	1 minute	Error	All	Check IMS reachability Try to reconnect IMS
Storage Discovery Host down	Checks if the discovery daemon is down on the storage discovery host	15 minutes	Error	Migrate	Resolve the discovery daemon issue

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
DNS removed	Checks if DNS is removed from the resiliency group where DNS customization is enabled	real time	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Edit the Resiliency Group and disable DNS customization
IOTap driver not configured	Checks if the IOTap driver is not configured	2 hours	Error	None	Configure the IOTap driver This risk is removed when the workload is configured for disaster recovery.
VMware Discovery Host Down	Checks if the discovery daemon is down on the VMware Discovery Host	15 minutes	Error	Migrate	Resolve the discovery daemon issue
VM restart is pending	Checks if the virtual machine has not been restarted after add host operation	2 hours	Error	Create resiliency group	Restart the virtual machine after add host operation
New virtual machine added to replication storage	Checks if a virtual machine that is added to a Veritas Replication Set on a primary site, is not a part of the resiliency group	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Rehearsal 	Add the virtual machine to the resiliency group
Replication lag exceeding RPO	Checks if the replication lag exceeds the thresholds defined for the resiliency group. This risk affects the SLA for the services running on your production data center	5 minutes	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Check if the replication lag exceeds the RPO that is defined in the Service Objective
Replication state broken/critical	Checks if the replication is not working or is in a critical condition for each resiliency group	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Contact the enclosure vendor. In case of Resiliency Platform Data Mover, or raise a support case with Veritas

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Remote mount point already mounted	Checks if the mount point is not available for mounting on target site for any of the following reasons: <ul style="list-style-type: none"> ■ Mount point is already mounted ■ Mount point is being used by other assets 	<ul style="list-style-type: none"> ■ Native (ext3, ext4, NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours 	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Unmount the mount point that is already mounted or is being used by other assets Risk gets resolved after 30 minutes if a successful cleanup rehearsal, migrate, or recover operation performed and VMware vCenter gets refreshed within 30 minutes.
Disk utilization critical	Checks if at least 80% of the disk capacity is being utilized. The risk is generated for all the resiliency groups associated with that particular file system	<ul style="list-style-type: none"> ■ Native (ext3, ext4, NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours 	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Rehearsal 	Delete or move some files or uninstall some non-critical applications to free up some disk space
ESX not reachable	Checks if the ESX server is in a disconnected state	5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop operations ■ On secondary site: migrate or recover operations 	Resolve the ESX server connection issue

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
vCenter Server not reachable	Checks if the virtualization server is unreachable or if the password for the virtualization server has changed	5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop operations ■ On secondary site: migrate or recover operations 	Resolve the virtualization server connection issue In case of a password change, resolve the password issue
vCenterDown	1.The vCenter server is down or unreachable. 2. The vCenter server is down. Unable to establish secure connection with vCenter server as the SSL/TLS handshake has failed.	15 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop, migrate, rehearsal, local recover, resync operations. ■ On secondary site: migrate, resync, recover, cleanup rehearse operations. 	For pt 1. Check for any one of the following issues and resolve: <ol style="list-style-type: none"> 1 vCenter server is down, vSphere service is not working, or vCenter server port has been changed after vCenter server configuration. 2 If port has been changed after configuration then, perform edit vCenter server operation to resolve the issue. For pt 2. Install valid CA certificates of vCenter server in the Resiliency Platform and refresh the vCenter server configuration

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Insufficient compute resources on failover target	Checks if there are insufficient CPU resources on failover target in a virtual environment	6 hours	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Reduce the number of CPUs assigned to the virtual machines on the primary site to match the available CPU resources on failover target
Host not added on recovery data center	Checks if the host is not added to the IMS on the recovery data center	30 minutes	Error	Migrate	Check the following and fix: <ul style="list-style-type: none"> ■ Host is up on recovery data center ■ Host is accessible from recovery datacenter IMS ■ Time is synchronized between host and recovery datacenter IMS
NetBackup Notification channel disconnected	Checks for NetBackup Notification channel connection state	5 minutes	Error	Recover	Check if the NetBackup Notification channel is added to the NetBackup primary server If the risk resolution or description indicates an SSL/TLS verification error. Refer this troubleshooting guide.
Backup image violates the defined RPO	Checks if the backup image violates the defined RPO	30 minutes	Warning	No operation	<ul style="list-style-type: none"> ■ Check the connection state of NetBackup Notification channel ■ Check for issues due to which backup images are not available

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
NetBackup primary server disconnected	Checks if NetBackup primary server is disconnected or not reachable	5 minutes	Error	Recover	Check if IMS is added as an additional server to the NetBackup primary server
NetBackup Recovery Host decommissioned	Check if NetBackup Recovery Host is disconnected or not reachable.	5 minutes	Error	recover	<ul style="list-style-type: none"> ■ Edit the resiliency group and choose different recovery host. ■ Try to connect the same recovery host.
Assets do not have copy policy	Checks if the assets do not have a copy policy	3 hours	Warning	No operation	Set up copy policy and then refresh the NetBackup primary server
Target replication is not configured	Checks if the target replication is not configured	3 hours	Warning	No operation	Configure target replication and then refresh the NetBackup primary server
Disabled NetBackup Policy	Checks if NetBackup policy associated with the virtual machine is disabled	3 hours	Warning	No operation	Fix the disabled policy

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Replication block tracking disk not found	Checks for the replication block tracking disk. If the replication block tracking disk is not found, then virtual machine does not get configured for remote recovery and the replication stops	30 minutes	Error	Migrate	<p>Ensure that the RBT disk is attached to the virtual machine. After the risk gets resolved, perform reboot of VM then perform the resync operation to avoid disk corruption during migrate or migrate back.</p> <p>If you are not able to locate the RBT disk then perform following steps in the order listed:</p> <ol style="list-style-type: none"> 1 Remove the virtual machine from resiliency group. 2 Add it again to a resiliency group to ensure that virtual machine is protected.
Members are manually deleted from network groups	Network group goes into faulted state when a member is manually removed. The risk is circulated to resiliency group	Immediate	Warning	Migrate, Rehearse	Edit the network group by adding the missing member and then edit the resiliency group details
Members deleted from network groups	Network group goes into faulted state when a discovered member gets deleted from IMS. The risk is circulated to resiliency group	5 minutes	Warning	Migrate, Rehearse	Edit the network group by adding the missing member and then edit the resiliency group details

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Virtual machine configuration not backed up	Unable to take a backup of virtual machine configuration file.	Immediate	Error	<ul style="list-style-type: none"> ■ Create resiliency group ■ Migrate ■ Rehearse 	Check the state of the IMS and its corresponding assets such as the hypervisors and vCenter servers. Perform edit resiliency group operation.
Unable to backup latest Virtual machine configuration	Unable to take a backup of the latest configuration file of the virtual machine.	Immediate	Warning	<ul style="list-style-type: none"> ■ Edit resiliency group ■ Migrate ■ Rehearse 	Check the state of the IMS and its corresponding assets such as the hypervisors and vCenter servers. Perform edit resiliency group operation.
Datastore for disk has changed to X, this datastore is not part of resiliency group	If virtual disk is moved to a non-compliant datastore. Applicable for 3rd party replication technology	5 to 15 minutes	Error	All operations except start and stop resiliency group	Edit the resiliency group or move the disk to a datastore which is part of the resiliency group.
Datastore for configuration file has changed to X, this datastore is not part of resiliency group. Previous datastore was Y.	If the virtual machine configuration file is moved to a non-compliant datastore. Applicable for 3rd party replication technology	5 to 15 minutes	Error	All operations except start and stop resiliency group	Edit the resiliency group or move the disk to a datastore which is part of the resiliency group.
Disk path has changed	Displayed when virtual machine snapshot is taken. Risk is resolved automatically after updating the blob.	5 to 15 minutes	Error	All operations	Risk is automatically resolved.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
New datastore added to the consistency group is not part of resiliency group	New datastore added to consistency group Applicable for 3rd party replication technology	6 hours	Error	<ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync 	Edit the resiliency group
Datastore removed from resiliency group	Datastore removed from consistency group Applicable for 3rd party replication technology	6 hours	Error	<ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync 	Edit the resiliency group
Veritas Replication VIB upgrade pending	Checks if the Veritas Replication VIB version on ESXi cluster has latest version installed.	6 hours	Error	None	Upgrade the Veritas Replication VIB to the latest version.
Veritas Replication VIB is in partial state.	Checks if the Veritas Replication VIB installation on ESXi cluster is in partial or unknown state.	6 hours	Error	<ul style="list-style-type: none"> ■ If the risk is on the target ESXi cluster then block the migrate and rehearsal operations. ■ If the risk is on the source ESXi cluster then block the resync operation. 	Perform Resolve and Verify operation on the ESXi cluster to fix the installation issues.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Insufficient privileges on vCenter server	Operations on the resiliency group may fail because of missing privileges on vCenter server data centers.	6 hours	Warning	One or more operations on resiliency group may fail because of missing privileges on vCenter server data center.	Ensure that appropriate privileges are configured on vCenter server data center before invoking any operation. Refer to the documentation for the required privileges.
Infrastructure Management Server data reporting disabled	Infrastructure Management Server cannot report data to Resiliency Manager due to version incompatibility	As soon as IMS connects to the Resiliency Manager after the Resiliency Manager upgrade	Error	All	Upgrade IMS to the latest version that is specified in the risk message
DRS Datastore Is Added Or Removed	New datastore is added to the cluster or is removed from the cluster	6 Hours	Warning	None	Edit the resiliency group
Datastore Cluster Deleted	Datastore cluster is deleted from the data center	6 Hours	Error	<ul style="list-style-type: none"> ■ Rehearsal ■ Migrate ■ Resync 	Edit the resiliency group
All the hosts on the applications are not reachable	All the hosts for the application are not reachable	15 minutes	Error	None	Check the connectivity with the application hosts
Application host is disconnected due to change in MAC address	Application Host is in Disconnected state	15 minutes	Error	<ul style="list-style-type: none"> ■ Rehearsal ■ Migrate 	Retry Add Host operation

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Assets does not have copy policy	Assets does not have copy policy	When vrp_host unassociated with copy policy.	Warning	None	Check if any asset has no copy policy
Backup image violates the defined RPO	Checks if the backup image violates the defined RPO	Immediate	Warning		<ul style="list-style-type: none"> ■ Check the connection state of NetBackup Notification channel. ■ Check for issues due to which backup images are not available.
CPU Usage Critical	Available compute capacity on the recovery site may be inadequate for recovering this application. This risk affects the recoverability of the services running on your production data center.	6 hours	Warning	None	Reduce the number of CPUs assigned to the virtual machines on the primary site to match the available CPU resources on failover target
Incorrect .Net version is installed	The expected .NET version is not installed or it is not compatible with the PowerShell version	2 hours	Error	<ul style="list-style-type: none"> ■ On Primary site: migrate and recover operations 	Ensure that the .NET version is installed with its compatible PowerShell version. Refer to the HSCL for compatible versions of .NET and PowerShell.
Editing the resiliency group is required	Resiliency group needs an upgrade or perform Edit operation.	Immediate	Warning	None	Edit the resiliency group using the Edit Configuration intent. Ensure that the resiliency group is online on the source datacenter before performing the edit operation

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Evacuation plan for data center has been invalidated.	Evacuation plan for data center has been invalidated, due to adding , deleting or updating a resiliency group or a VBS	Immediate	Error	None	Regenerate the evacuation plan.
Host reboot is pending after upgrade	The OS is not rebooted after upgrade operation	Immediate	Warning	None	Virtual machine requires to be rebooted after the upgrade operation
Mount point is deleted	Check if the mount point on which the assets of the resiliency group are configured, is deleted or renamed	6 hours	Error	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal 	Remount using the same mount point else you need to edit the resiliency group
PowerShell is not initialized	PowerShell is not initialized	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and rehearsal operations 	Check PowerShell Initialization on host
PowerShell is not installed	PowerShell is not installed	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and rehearsal operations 	Install PowerShell (version > 2.0) on host
Powershell Version is incorrect	Expected Powershell version not found	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and recover operations 	Install Powershell version should be 2.0 and above

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Registry Parameter LSI_SAS is not set	Registry Parameter LSI_SAS is not set	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and rehearsal operations 	Change the value for registry parameter LSI_SAS->Start to 0 and refresh host discovery
Replication Gateway is not reachable	The Replication Gateway is down or not reachable from the IMS	15 minutes	Error	None	Make sure the replication gateway appliance is running and is reachable from the IMS
Virtual machine or Replication Gateway is not found in the cluster.	Virtual machine or Replication Gateway is not present in cluster using which resiliency group is created.	15 minutes	Error	migrate, rehearsal	<ul style="list-style-type: none"> ■ If the virtual machine is moved out of the cluster or ESX server, re-add the virtual machine and perform resync operation. ■ If the virtual machine is unregistered, perform start resiliency group operation on the resiliency group with the checkbox "Refresh storage, network, compute, and customization" as selected.
Replication state synchronizing	Data synchronization is in progress.	5 minutes	Warning	None	Wait for synchronization to complete (Replication state should be Active (Connected Consistent))

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Resync operation is pending on a resiliency group	Resync operation is pending on current resiliency group	Immediate	Error	On Secondary site: migrate operation	Execute Resync operation on current resiliency group
Resiliency group configuration drift	Disk configuration for asset(s) in the resiliency group is changed. This is a configuration drift.	2 minutes	Error	<ul style="list-style-type: none"> ■ On Primary site: rehearsal operation ■ On Secondary site: migrate, resync, and rehearsal operations 	Refresh the respective hosts, vCenter servers or Hyper-V servers and the cloud discovery. After refresh, probe the risk. If the risk still exists, remove the virtual machine from the resiliency group and re-add using the edit operation.
Resiliency group configuration error	The disk size of the virtual machine in the resiliency group has changed. This is a configuration error	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and resync operation 	<ul style="list-style-type: none"> ■ Editing the size of a disk is not supported. Restore the disk size for a resiliency group having multiple virtual machines. ■ Edit the resiliency group by removing affected hosts and then add it again to re-protect. ■ For resiliency group having only one virtual machines delete it and recreate again.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Resiliency group outage in datacenter	Outage has been declared for the resiliency group in the datacenter	Immediate	Error	None	Perform remediation steps to clear outage in the specified datacenter. Run a Resync or Clear outage operation (as applicable) to indicate that the outage has been cleared
Data sync failed between Resiliency Manager and database.	Data sync failed between Resiliency Manager and database.	As soon as the vrp_rm vertex gets updated with property db_status as value "Data sync failed"	Error	None	Perform Resync operation for Resiliency Manager
SAN Policy Offline Shared	SAN policy on the Windows host is Offline Shared	2 hours	Warning	None	Change the SAN policy on the Windows host to Online Shared and refresh the host discovery information

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Stale configuration :: Object deleted	Asset is unavailable	As soon as discovery reports delete of addressable objects.	Error	<ul style="list-style-type: none"> ■ On Primary site: Start, stop, migrate, resync, recover, and cleanup rehearsal operations. ■ On Secondary site: Migrate, rehearsal, resync, and recover operations. 	Edit the resiliency group.
Stale configuration :: Object unreachable	Asset is unreachable	As soon as discovery reports DISCONNECTED or NOT REACHABLE fault for addressable objects.	Error	<ul style="list-style-type: none"> ■ On Primary site: Start, stop, migrate, , resync , recover, and cleanup operations. ■ On Secondary site: Migrate, rehearsal resync, and recover operations. 	Edit the resiliency group.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
The migrated virtual machine is not added to the target IMS.	The migrated virtual machine is not added to the target IMS.	45 minutes	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and resync operation 	Refer to the documentation to know the possible reasons for failure of add host operation
Unable to get VMX	Unable to backup virtual machine configuration file	Immediate	Error	<ul style="list-style-type: none"> ■ On Primary site: rehearsal, migrate and recover operations 	Check the state of IMS, its corresponding assets such as the hypervisors and vCenter servers. Perform edit resiliency group operation.
Unable to update virtual machine configurations file	Unable to backup latest virtual machine configuration	Immediate	Error	None	Check the state of IMS, its corresponding assets such as the hypervisors and vCenter servers. Perform edit resiliency group operation.
vCenter server is removed from IMS	vCenter server is removed from IMS	Immediate	Error	<ul style="list-style-type: none"> ■ On Primary site: start, stop, rehearsal, and migrate operations ■ On Secondary site: start, stop, rehearsal, and migrate operations 	Add the vCenter server to the IMS.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
VCS Servicegroup Faulted	VCS Servicegroup is in Faulted state	1 hour	Error	None	Resolve the fault on VCS Servicegroup
Insufficient quota on target vCloud Director	Sufficient quota(CPUs/Memory/Storage) is not available on target vCloud Director.	5 minutes	Error	None	Sufficient quota should be available on the target vCloud Director
Virtual machine is deleted	One or more virtual machines are deleted or unregistered. The virtual machines belong to a resiliency group that is configured for remote recovery. This affects the recoverability of the resiliency group.	6 hours	Error	On Secondary site: migrate operation	Edit the resiliency group to remove the virtual machines that are deleted or unregistered.
Virtual machine is not protected	Virtual machine is not configured for remote recovery	Immediate	Error	None	If the virtual machine is in production data center then configure the virtual machine for remote recovery. If the virtual machine is in vCloud data center then ensure that disk.EnableUUID property is set to TRUE on the VRP_VAPP_TEMPLATE virtual machine as well as on the migrated virtual machine. After the risk is resolved, perform the Resync operation to avoid disk corruption during migrate or migrate back operation.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
VMware discovery failed	VMware discovery is failed . Unable to establish secure connection with vCenter server as the SSL/TLS handshake has failed	6 hours	Error	None	1. In case of a password change, resolve the password issue and refresh the vCenter server configuration. 2. Install valid CA certificates of vCenter server in the Resiliency Platform and refresh the vCenter server configuration.
IO Filter is not replicating the IOs from the virtual machine	IO Filter has encountered a fatal error	When IMS is receiving NOOP snmp event.	Error	<ul style="list-style-type: none"> ■ On Primary site: migrate resync deepstart (Perform start operation after reverse replication is complete) ■ On Secondary site: migrate resync deepstart (Perform start operation after reverse replication is complete) 	If IO filter has encountered errors, either invoke the edit resiliency group workflow to remove and re-add asset from the resiliency group or delete the resiliency group and create it again

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Cloud discovery failed	<p>1. Cloud discovery has failed.</p> <p>2. Unable to establish secure connection with cloud resource as the SSL/TLS handshake has failed.</p>	<p>1. After 5 minutes</p> <p>2. After 5 minutes</p>	<p>1. Error</p> <p>2. Error</p>	<ul style="list-style-type: none"> ■ On primary site: migrate, recover, rehearsal, cleanup rehearsal, and resync operations. ■ On secondary site: start, stop, migrate, and resync operations 	<p>1. Edit the cloud configuration to resolve the issue. If risk persists contact Veritas Support.</p> <p>2. Install valid CA certificate of the cloud resource in Resiliency Manager and then refresh the cloud configuration.</p>
Enclosure failed	Enclosure discovery has failed.	After 5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: migrate, recover, rehearsal, cleanup rehearsal, and resync operations. ■ On secondary site: migrate, and resync operations 	<p>General: Edit the enclosure configuration and provide valid SSH host keys. If the risk persists, contact Veritas Support.</p> <p>NetApp SSL resolution: Install valid CA certificates of NetApp enclosure in the Resiliency Platform and refresh the NetApp enclosure configuration.</p> <p>NetApp general resolution: Unable to fetch enclosure details. If the risk persists contact Veritas Support.</p>

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Cloud authentication failed	Cloud credentials are incorrect	After 5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: migrate, recover, rehearsal, cleanup rehearsal, and resync operations. ■ On secondary site: start, stop, migrate, and resync operations 	Edit cloud configuration and provide correct credentials to resolve the issue. In case of AWS, check the IAM role with proper privileges is attached to IMS.
Cloud connection timeout	Connection timed out fetching information about cloud resources.	After 5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: migrate, recover, rehearsal, cleanup rehearsal, and resync operations. ■ On secondary site: start, stop, migrate, and resync operations 	Resolve network connectivity between IMS and cloud data center and then refresh the cloud configuration.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
NTP Time Sync Failed	NTP time skew. Time skew must be less than 3 seconds.	5 minutes	Warning	None	Synchronize with NTP server.
NTP Time Unsynchronized	Not able to synchronize with the NTP server.	5 minutes	Warning	None	Synchronize with NTP server.
NTP Time Indeterminate	NTP status indeterminate	5 minutes	Warning	None	Synchronize with NTP server.
Resiliency Group Configuration Drift for Network changed of some of the assets in the Resiliency Group	This risk is raised if network of some of the assets in the Resiliency Group is changed after the Resiliency Group is created. The change can be in the VLAN, vSwitch or cloud network settings.		Error		The risk is resolved when the deleted network gets discovered in Veritas Resiliency Platform. Or the network update risk will be resolved after successful editing the Resiliency Group.
The network setting of the virtual machine 'Connect at Power On' is disabled.	The network setting of the virtual machine 'Connect at Power On' is disabled.	Immediate	Warning	None	Ensure to set the 'Connect At Power On' network adapter settings of the virtual machine is enabled.
Node added to Infoscale Cluster	A new node is added to InfoScale cluster and the node has not been configured into Resiliency Platform yet.	Raised when Resiliency Platform InfoScale cluster discovery detects the node addition in near real time after the node has been added to the cluster.	Warning	None	Perform reconfigure operation on the InfoScale cluster to configure the node in Resiliency Manager.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Appliance Storage critical	At least 85% of the disk capacity is being utilized. This risk affects the continuity of the services running on your appliance.	5 min	Warning	All	Free up the disk space or increase the disk size of the appliance.
ISO already mounted	ISO is mounted on the host	'5 min	Error	Rehearsal, Cleanup rehearsal	Remove or detach the ISO from VMware virtual machine.
Replication Gateway Service Down	One or more services on the Replication Gateway appliance are not running.	5 min	Error	<ul style="list-style-type: none"> ■ On the primary site: migrate and resync operations ■ On secondary site: migrate, recover , resync , rehearsal and, cleanup rehearsal operations. 	Make sure all the services on the Replication Gateway appliance are running.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Replication state is inactive	This risk affects the recoverability of the services running on your source data center.	10 mins	Warning	None	

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
					<p>Perform the following steps:</p> <ol style="list-style-type: none"> 1 Make sure the source Replication Gateway is powered on and all services are running. 2 In case of in-guest replication, if risk is raised: <ul style="list-style-type: none"> ■ Make sure workload is powered on.. ■ Make sure workload should communicate with source Replication Gateway over the network 3 In case of VAIO: <ul style="list-style-type: none"> ■ Make sure ESXi server on which virtual machine is residing is able to communicate with source Replication Gateway over the network. ■ In case target Replication Gateway is replaced before migrate operation and

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
					<p>after migrate operation, risk is raised on the resiliency group, perform a resync operation for resiliency group. The resync operation in this case will perform diff sync. If resync is performed on resiliency group after migrate operation, usually it performs full sync. Hence, resync needs to be done carefully for those resiliency group which has this risk after successful migrate operation.</p> <ul style="list-style-type: none"> ■ In case vMotion happens for virtual machine is being snapshotted, this risk may appear. but it gets resolved in next CG update.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
					<ul style="list-style-type: none"> ■ If a policy is detached from the workload disk, a risk appears. Performing resync operation should resolve the risk.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
					<p>4 If risk is raised on Replication Gateway after VBS migrate operation.</p> <ul style="list-style-type: none"> ■ If first attempt of VBS migrate operation fails in reverse-replication step after updating CG roles, consecutive VBS migrate operation will skip reverse replication and replication will remain inactive. Here, check source CG state on gateway. If the CG state is set to “stopped”, then perform resync operation for individual resiliency group. If resync is performed on VBS, full sync may occur for some of the resiliency group.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Replication configuration is broken for resiliency group	Replication configuration is broken for resiliency group.	Real time after recovering the resiliency groups configured with multiple recovery point.	Error	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal ■ Recover with replication path ■ Cleanup rehearsal ■ Start operation with network and storage refresh 	Perform resync operation to repair the replication for the recovered virtual machine.
Protection configuration data for replication configuration is being updated for the resiliency group.	The protection configuration data for replication configuration is getting updated for the resiliency group.	Real time after recovering the resiliency groups configured with multiple recovery point.	Error	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal ■ Recover with replication path ■ Cleanup rehearsal ■ Start operation with network and storage refresh 	Once the replication configuration is updated successfully, this risk goes away.

Table 1-72 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
VMWare ESXi Cluster Membership Change	Checks if any new ESX host is added to the ESX cluster.	30 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal 	To resolve the risk perform the below steps :- 1 Create or change the required network configurations according to the environment and then Edit the Resiliency Group with Edit Configuration intent. 2 Remove the ESXi host from the cluster.
Maintenance Mode enabled on VMware ESXi host	Checks if maintenance mode is enabled on any of the ESX hosts.	30 minutes	Error / Warning	<ul style="list-style-type: none"> ■ Migrate ■ Takeover ■ Rehearsal ■ Cleanup Rehearsal 	Disable the Maintenance Mode on the ESX host or remove the ESXi host from the cluster.

See [“About risks”](#) on page 542.

About reports

Using the Veritas Resiliency Platform console, you can generate a variety of reports. The following are the broad categories under which the reports are grouped:

- **Inventory:** Reports in this category provide information about the data centers and applications, and the virtual machines that are deployed in the data centers.
- **Recovery Assessment:** This category lists the reports that are related to the disaster recovery operations such as the migrate and take over report, and the rehearsal report.

- **Risk:** This category has two reports; Current Risk and Risk History. These reports show the summary and details of all the current and historical risks that occurred in the environment.

Reports can be scoped on the data center or global. You can subscribe for a report on a daily, weekly, monthly, quarterly, or yearly basis, or on predefined days of the week, or run on demand. Reports are available in the HTML and PDF format, or as a comma-separated file (CSV) file.

You can send a report to multiple recipients by entering the email addresses separated by a comma (,) or a semicolon (;).

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 (navigation pane)

Click **Inventory**, **Recovery Assessment**, or **Risk** to expand the category.

- 2 Click **Schedule** on the desired report.
- 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 (;).

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 (navigation pane)

Click **Inventory**, **Recovery Assessment**, or **Risk** to expand the category.

2 Click **Run** on the desired report.

3 In the **Run Report** wizard panel, specify the following information, 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 (;).

Viewing reports

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

Resiliency Groups and VBS Summary	Provides details about the resiliency groups and VBSs in the data centers across all sites.
-----------------------------------	---

VM Inventory	<p>Provides the platform distribution and the OS distribution details of the virtual machines that are deployed in the data centers in the form of a pie chart.</p> <p>The Details table provides additional information for each virtual machine.</p> <p>For virtual machines on the Hyper-V Server, the report displays the total memory instead of allocated memory.</p> <p>Hyper-V virtual machines which are in offline state are displayed in the Unknown category.</p>
License Entitlement Report	<p>License Entitlement Report provides details about the licenses that are deployed in your datacenter. There are 3 types of licenses:</p> <ul style="list-style-type: none"> ■ Veritas Resiliency Platform FETB (Per-FETB) ■ Veritas Resiliency Platform Compute (Per-Core) ■ Veritas Resiliency Platform Compute (Per-VM)
Notification Throttling Report	<p>Notification Throttling Report displays all the notifications which are currently throttled and are waiting to be raised.</p>
Activity Distribution History	<p>Provides information about tasks, such as migrate, recover, rehearse, start, and stop, performed for a specified duration.</p>
Recovery Activity History by RG	<p>Provides historical information about recovery tasks, such as migrate, recover, and rehearse for each resiliency group.</p>
Recovery Activity History by VBS	<p>Provides historical information about recovery tasks, such as migrate, recover, and rehearse for each VBS.</p>
Metering	<p>Provides details of the virtualization servers that are protected for disaster recovery.</p> <p>You can view the total number of servers that are protected for disaster recovery. For these servers you can view the total memory, processor cores, and the total storage.</p>

VBS RPO

Provides Recovery Point Objective (RPO) details for all the virtual business services (VBS) in the resiliency domain.

The bar chart provides information on the top VBS with maximum RPO lag.

You can view the lag in the last replication and the replication date for all the VBS in the table.

To view a report

1 Navigate



Reports (navigation pane)

Click **Inventory**, **Recovery Assessment**, or **Risk** to expand the category.

2 Do one of the following:

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

Managing a running activity

Using the Veritas Resiliency Platform console, you can abort a task or an operation which is currently running. And you can also resume an operation which is in Pause state.

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

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

While configuring a resiliency group or a Virtual Business Service (VBS) for disaster recovery, you can select the pause or the manual intervention points. These manual intervention points let you pause the Migrate and Recover operations. You can

resume the workflow by selecting the Resume option on **Current** activities page or **Recent Activities** page.

See “[About manual intervention](#)” on page 525.

To abort an activity

1 Navigate

Do one of the following:



Activities (navigation pane). Skip to [2](#)

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

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

- Right click and select **Abort**.
- Click on the vertical ellipsis and select **Abort**
- Right click and select **Details**. Click **Abort** on the details page.

To resume a paused activity

1 Navigate

Do one of the following:



Activities (navigation pane). Skip to [2](#)

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

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

- Right click and select **Resume**.
- Click on the vertical ellipsis and select **Resume**
- Right click and select **Details**. Click **Resume** on the details page.

Miscellaneous references

After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish

menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.

- See [“About klish”](#) on page 584.
- See [“About applying updates to Resiliency Platform”](#) on page 631.
- See [“Virtual appliance security features”](#) on page 667.

About klish

Once the Veritas Resiliency Platform virtual appliance is deployed and configured, you are given limited, menu-based access to the operating system and the product. You need to use Command Line Interface Shell (klish) menu to manage the configuration-related changes to the product.

Below are the Klish options:

Table 1-73 Klish main menu options

Menu option	Description
manage	Manage the Veritas Resiliency Platform appliance
monitor	Monitor the Veritas Resiliency Platform appliance activities
network	Change some of the network configurations
settings	Change the system settings
hotfix	Manage the Veritas Resiliency Platform hotfixes
support	Access the Veritas Resiliency Platform logs
updates	Manage Veritas Resiliency Platform updates and patches
utilities	Run the miscellaneous utilities of the appliances

After the product configuration, whenever you log in to the Resiliency Platform appliance, you get the main menu of klish. This menu is the starting point, from which you can configure, manage, monitor, and support your application using the command line. You can reconfigure or modify some of the appliance settings that are configured through the product bootstrap. Following are the settings that you can reconfigure using klish:

- **Network settings:**
You can reconfigure the subnet mask, IP, default gateway, DNS server, route, traceroute, SSH-enable-NIC, NIC for accessing product user interface and search domains using the klish menu. You cannot reconfigure the hostname

that you had configured through the bootstrap process. In case of static DHCP, you cannot change the network settings using the klish menu. You cannot change the network settings for any component that is configured in the cloud environment.

- **System settings:**

You can reset the time zone, perform operations related to NTP server, shut down the appliance, reboot using the klish menu. Changing the system settings can affect the product functionality if incorrect values are set. You can also perform logical volume management (LVM) operations such as adding a disk or removing a disk using the klish menu.

- **About updates:**

You can apply patch updates on Resiliency Platform virtual appliance for update, rollback a previously prepared update and view the latest version of the Resiliency Platform. You can also configure the repository, display the current repository configuration and remove the repository.

How to use help in Klish

You can press the **tab** key to display the menu options or you can run the `help` command to get detailed help on how to use klish. Use **space** key for auto-completion of command. If you get the `Syntax Error: Illegal command line error` or `Syntax Error: The command is not completed error`, press **?** key to display detailed help on the required parameter.

Lock mode in Klish

If a klish command is expected to perform any operation on an entity such as start or stop services, it goes into lock mode and does not allow any other operation from any other session to be performed till the first operation gets completed. In such a scenario, you may encounter the following warning:

```
Warning: Try to get lock. Please wait...
```

After waiting for some time, if the operation still cannot be performed due to the lock, then you may encounter the following error:

```
Error: Can't get lock
```

In this case, you need to execute the same command after waiting for some time. The operation is performed if the lock gets released by that time.

Best practice for using Klish

At times, you may not be able to run the klish commands if the `/var/opt` directory is fully utilized and there is no space to run the klish commands. We now raise risks if disk space is running low on any appliance. Once this issue occurs, there is no way to recover from this situation. Hence, you need to periodically check if the

space in that directory is getting fully occupied, and provision for an extra disk accordingly.

See “[Klish menu options for Resiliency Manager](#)” on page 586.

See “[Klish menu options for IMS](#)” on page 598.

See “[Klish menu options for Replication Gateway](#)” on page 612.

Klish menu options for Resiliency Manager

Following are the options available for Resiliency Manager using klish menu:

Table 1-74 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
hotfix	Manage hotfixes Table 1-75
manage	Manage appliance Table 1-76
monitor	Monitor appliance activities Table 1-78
network	Manage network configuration Table 1-79
settings	Manage appliance settings Table 1-88
support	To access logs Table 1-97
updates	Manage updates and patches Table 1-101
utilities	Run miscellaneous utilities Table 1-102

Table 1-75 Options available with **hotfix** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
apply-hotfix	Apply the specified hotfix
list-applied-hotfixes	List the applied hotfixes
list-available-hotfixes	List the available hotfixes
uninstall-hotfix	Uninstall the specified hotfix

Table 1-76 Options available with **manage** command

Menu option	Description
back	Return to the previous menu
configure	Configure Resiliency Platform component or show the configured component
exit	Log out from the current CLI session
services	Manage the appliance services Table 1-77
help	Display an overview of the CLI syntax

Table 1-77 Options available with **services** command

Menu option	Description
force	Perform operations forcefully by skipping services validations. <ul style="list-style-type: none">■ force restart service name command restarts the service name mentioned forcefully.■ force stop service name command stops the service name mentioned forcefully. You can provide multiple service names (comma separated) or can provide ALL for all services

Table 1-77 Options available with **services** command (*continued*)

Menu option	Description
restart	Restart Resiliency Platform services Two options available are: <code>restart all</code> where, <i>all</i> means all the services. <code>restart service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).
start	Start Resiliency Platform services Two options available are: <code>start all</code> where, <i>all</i> means all the services. <code>start service name</code> where, <i>service name</i> 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: <code>status all</code> where, <i>all</i> means all the services. <code>status service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).
stop	Stop Resiliency Platform services Two options available are: <code>stop all</code> where, <i>all</i> means all the services. <code>stop service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).

Table 1-78 Options available with **monitor** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
top	Display the top process information

Table 1-78 Options available with **monitor** command (*continued*)

Menu option	Description
who	Display who is currently logged into the appliance
uptime	Display the uptime statistics for the appliance
FSuage	Display filesystem usage

Table 1-79 Options available with **network** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
dns	Show or set the DNS server or manage the options for resolv.conf file Table 1-80
ip	Show the IP address Table 1-82
route	View and manipulate the IP routing table Table 1-83
search-domain	Show or change the domain Table 1-84
traceroute	Trace packet routes to a particular host. You can also specify a port to trace the packet routes.
ssh-enabled-nic	Show or update SSH enabled NIC Table 1-85
nic-configuration	Show and configure the NIC Table 1-86
nic-for-UI	Show or update NICs configured to access product user interface Table 1-87

Table 1-80 Options available with **dns** command

Menu option	Description
options	Show, add, or remove options to the <code>/etc/resolv.conf</code> file. Refer to the documentation of <code>resolv.conf</code> for a list of available options and their purpose. Table 1-81
set	Configure Domain Name Server
show	Show the current Domain Name Server

Table 1-81 Options available with **options** command

Menu option	Description
add	Add a <code>resolv.conf</code> option
remove	Remove a <code>resolv.conf</code> option
show	Show options of <code>resolv.conf</code> file

Table 1-82 Options available with **IP** command

Menu option	Description
show	Show the current IP address

Table 1-83 Options available with **route** command

Menu option	Description
add	Set a default route or a route for a host or a subnet
delete	Delete the route entry from the routing table
show	Display your current routing table

Table 1-84 Options available with **search-domain** command

Menu option	Description
add	Add a search-domain
remove	Remove the search domain name
show	Show the search domain settings

Table 1-85 Options available with **ssh-enabled-nic** command

Menu option	Description
show	Show the NICs on which SSH is enabled. By default, SSH is enabled on all the NICs.
add	Add network interface to enable SSH on it
remove	Remove network interface to disable SSH on it

Table 1-86 Options available with **nic-configuration** command

Menu option	Description
show	Show details of NIC configuration like hostname, IPv4 or IPv6 address, prefix, gateway etc.
set	Configure the NICs which are not used while bootstrapping.

Table 1-87 Options available with **nic-for-UI** command

Menu option	Description
show	Show the NICs which are used to access product web user interface.
set	Set a NIC to access product web user interface from existing configured NICs.
remove	Remove one of the NICs used to access product web user interface.

Table 1-88 Options available with **settings** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
date	Display the current date and time for the appliance Table 1-89
lvm	Perform operations related to logical volume manager on the Appliance Table 1-90

Table 1-88 Options available with **settings** command (*continued*)

Menu option	Description
ntp	Perform operations related to NTP server
change-password	Change the admin user password for the appliance
poweroff	Shut down the appliance
reboot	Restart the appliance
timezone	Show or change the timezone for the appliance Table 1-94
password-policies	Perform operation related to password policies of administrator user for the appliance. Table 1-95

Table 1-89 Options available with **date** command

Menu option	Description
show	Show the time and date

Table 1-90 Options available with **lv** command

Menu option	Description
add-disk	Add disk to the OS or data volume. You need to attach a disk before adding it. Table 1-91
list-free-disk	List the free disks
initialize-free-disk	Initialize the newly attached free disk
list-used-disk	List the disks used by the OS or data volume Table 1-92

Table 1-90 Options available with **lvm** command (*continued*)

Menu option	Description
remove-disk	<p>Remove disk from the data volume. Remove disk operation involves migrating data from the existing disk to a new disk. You can remove a disk only after attaching a new disk with enough storage to migrate the data.</p> <p>The command first displays the list of disks being used and you need to select the disk that you want to remove. Then it displays the list of free disks where you want to migrate data and you need to select the disk. You can choose whether to initialize the new disk or not.</p> <p>It is recommended to suspend replication of all the configured Veritas Replication Sets before performing the remove disk operation.</p>
resize-logicalvolume	<p>Resize the OS or data volume for the resized data disk</p> <p>Table 1-93</p>

Note: In case you initialize the newly-added disk during add-disk or remove-disk operation, the existing data on the new disk is deleted.

Table 1-91 Options available **add-disk** command

Menu Options	Description
data-volume	Add disk to the data volume
os-volume	Add disk to the OS volume

Table 1-92 Options available with **list-used-disk** command

Menu Options	Description
data-volume	Lists disks used by the data volume
os-volume	Lists disks used by the OS volume

Table 1-93 Options available with **resize-logical-volume** command

Menu Options	Description
data-volume	Resize the data volume
os-volume	Resize the OS volume

Table 1-94 Options available with **timezone** command

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

Table 1-95 Options available with **password-policies** command

Menu option	Description
set	Modify the administrator user password policies for the appliance. Table 1-96
show	Show the administrator user password-policies for the appliance.

Table 1-96 Options available with **password-policies set** command

Menu option	Description
max-age	Modify the maximum number of days before password change is required for administrator user.
min-age	Modify the minimum number of days before password change is required for administrator user.
min-length	Modify minimum password length for the administrator user.
warning-days	Modify number of days before a warning for administrator password expiry is given.

Table 1-97 Options available with **support** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
reset-support-password	Reset the support user password to the default installation password. This option may typically be used by Veritas support.

Table 1-97 Options available with **support** command (*continued*)

Menu option	Description
shell	Open the bash shell prompt for support user
loggather	If the appliance is configured as a Resiliency Manager, then various options will be available for collecting the Resiliency Manager logs. Table 1-98

Table 1-98 Options available with **loggather** command

Menu option	Description
basic	Gather logs of Resiliency Manager excluding database and heap dumps See Table 1-99 on page 595.
full	Gather logs of Resiliency Manager with database
db	Gather database logs of Resiliency Manager
coredump	Gather heap dumps of Resiliency Manager See Table 1-100 on page 596.
cleanup coredump	Clean up all the collected loggater heap dump files of Resiliency Manager
cleanup vrp-logs	Clean up all the collected loggater log files of Resiliency Manager
cleanup all	Clean up all the collected loggater files (vrp-logs and coredump) of Resiliency Manager
show	Lists all the loggater URLs ordered by date and time of Resiliency Manager

Table 1-99 Options available with **basic** command

Menu option	Description
Number of days	Displays the basic logs from the days (1-99) mentioned

Table 1-100 Options available with **coredump** command

Menu option	Description
days	Displays the coredump logs from the days (1-99) mentioned

Table 1-101 Options available with **updates** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
Show-version	Show the appliance version
prepare-for-update	Save the virtual appliance configuration in preparation for upgrade
rollback-update	To rollback the prepare for update operation

Table 1-102 Options available with **utilities** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
clear	Clear the screen
unmount-cd-rom	Unmount the CD-ROM from the appliance
troubleshoot run-tool	Use the troubleshooting menu options Table 1-103
vmware-tools	Perform VMware Tools operations (install, uninstall, and show-version) Table 1-104
sftp-session	Use SFTP session for file transfer operation on the SFTP server Table 1-105

Table 1-102 Options available with **utilities** command (*continued*)

Menu option	Description
azure-waagent-service	Perform Azure waagent service operation. Applicable only in Azure environment Table 1-109

Table 1-103 Options available with **troubleshoot run-tool** command

Menu option	Description
view-logs	View log files on any virtual appliance
check-port	Verify required open ports on Veritas Resiliency Platform VSA for communication with other appliance using admin password.

Table 1-104 Options available with **vmware-tools** command

Menu option	Description
install	Install the VMware Tools mounted on CD-ROM of the appliance
show-version	Show the installed version of VMware Tools on the appliance
uninstall	Uninstall the VMware Tools from the appliance

Table 1-105 Options available with **sftp-session** command

Menu option	Description
start	To start the SFTP server session Table 1-106
show-details	View the current SFTP user and session details
stop	To stop the SFTP server session

Table 1-106 Options available with **start** command

Menu option	Description
get	View the file types that can be downloaded from the SFTP server Table 1-107

Table 1-106 Options available with start command (*continued*)

Menu option	Description
put	View the file types that can be uploaded on the SFTP server Table 1-108

Table 1-107 Options available with **get** command

Menu option	Description
logs	Download the log files and directories from the SFTP server
heap-dump	Download the heap dump files of the service available on Resiliency Manager

Table 1-108 Options available with **put** command

Menu option	Description
patch	Upload the private patch on the SFTP server

Table 1-109 Options available with **azure-waagent-service** command

Menu option	Description
start	Start Azure waagent service
stop	Stop Azure waagent service
status	Show current status of Azure waagent service

Klish menu options for IMS

Following are the options available for IMS using klish menu:

Table 1-110 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

Table 1-110 Options available in the **main** menu (*continued*)

Menu option	Description
hotfix	Manage hotfixes Table 1-111
manage	Manage appliance Table 1-112
monitor	Monitor appliance activities Table 1-116
network	Manage network configuration Table 1-117
settings	Manage appliance settings Table 1-126
support	To access logs Table 1-135
updates	Manage updates and patches Table 1-138
utilities	Run miscellaneous utilities Table 1-140

Table 1-111 Options available with **hotfix** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
apply-hotfix	Apply the specified hotfix
list-applied-hotfixes	List the applied hotfixes
list-available-hotfixes	List the available hotfixes
uninstall-hotfix	Uninstall the specified hotfix

Table 1-112 Options available with **manage** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
configure	Configure Resiliency Platform component or show the configured component Table 1-113
infra-appliance	List or remove Replication Gateway appliance Table 1-114
services	Manage the appliance services Table 1-115
show-resiliency-domain	Show details of resiliency domain to which the Infrastructure Management Server is configured. The option also lists the configured Resiliency Manager(s) in the domain.

Table 1-113 Options available with **configure** command

Menu option	Description
ims_register	Register the IMS using the registration URL obtained after initiating the Add IMS operation This option is available only for an IMS appliance Add link to Add an IMS topic
show	Show the configured component

Table 1-114 Options available with **infra-appliance** command

Menu option	Description
list	List Resiliency Platform infrastructure appliance.
remove	Remove the Replication Gateway appliance. You need to remove the Gateway pair before you remove the Gateway. Remove replication gateway link topic

Table 1-115 Options available with **services** command

Menu option	Description
restart	Restart Resiliency Platform services Two options available are: <code>restart all</code> where, <i>all</i> means all the services. <code>restart service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).
start	Start Resiliency Platform services Two options available are: <code>start all</code> where, <i>all</i> means all the services. <code>start service name</code> where, <i>service name</i> 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: <code>status all</code> where, <i>all</i> means all the services. <code>status service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).
stop	Stop Resiliency Platform services Two options available are: <code>stop all</code> where, <i>all</i> means all the services. <code>stop service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).

Table 1-116 Options available with **monitor** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
top	Display the top process information

Table 1-116 Options available with **monitor** command (*continued*)

Menu option	Description
who	Display who is currently logged into the appliance
uptime	Display the uptime statistics for the appliance
check-cim-status	Display the status of Common Information Model (CIM)
FSuage	Display filesystem usage

Table 1-117 Options available with **network** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
dns	Show or set the DNS server or manage the options for resolv.conf file Table 1-118
ip	Show the IP address Table 1-120
route	View and manipulate the IP routing table Table 1-121
search-domain	Show or change the domain Table 1-122
traceroute	Trace packet routes to a particular host. You can also specify a port to trace the packet routes.
ssh-enabled-nic	Show or update SSH enabled NIC Table 1-123
nic-configuration	Show and configure the NIC Table 1-124
nic-for-UI	Show or update NICs configured to access product user interface Table 1-125

Table 1-118 Options available with **dns** command

Menu option	Description
options	Show, add, or remove options to the <code>/etc/resolv.conf</code> file. Refer to the documentation of <code>resolv.conf</code> for a list of available options and their purpose. Table 1-119
set	Configure Domain Name Server
show	Show the current Domain Name Server

Table 1-119 Options available with **options** command

Menu option	Description
add	Add a <code>resolv.conf</code> option
remove	Remove a <code>resolv.conf</code> option
show	Show options of <code>resolv.conf</code> file

Table 1-120 Options available with **IP** command

Menu option	Description
show	Show the current IP address

Table 1-121 Options available with **route** command

Menu option	Description
add	Set a default route or a route for a host or a subnet
delete	Delete the route entry from the routing table
show	Display your current routing table

Table 1-122 Options available with **search-domain** command

Menu option	Description
add	Add a search-domain
remove	Remove the search domain name
show	Show the search domain settings

Table 1-123 Options available with **ssh-enabled-nic** command

Menu option	Description
show	Show the NICs on which SSH is enabled. By default, SSH is enabled on all the NICs
add	Add NIC to enable SSH on it
remove	Remove NIC to disable SSH on it

Table 1-124 Options available with **nic-configuration** command

Menu option	Description
show	Show details of NIC configuration like hostname, IPv4 or IPv6 address, prefix, gateway etc.
set	Configure the NICs which are not used while bootstrapping.

Table 1-125 Options available with **nic-for-UI** command

Menu option	Description
show	Show the NICs which are used to access product web user interface.
set	Set a NIC to access product web user interface from existing configured NICs.
remove	Remove one of the NICs used to access product web user interface.

Table 1-126 Options available with **settings** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
date	Display the current date and time for the appliance Table 1-127
lvm	Perform operations related to logical volume manager on the Appliance Table 1-128

Table 1-126 Options available with **settings** command (*continued*)

Menu option	Description
ntp	Perform operations related to NTP server
change-password	Change the admin user password for the appliance
poweroff	Shut down the appliance
reboot	Restart the appliance
timezone	Show or change the timezone for the appliance Table 1-132
password-policies	Perform operation related to password policies of administrator user for the appliance. Table 1-133

Table 1-127 Options available with **date** command

Menu option	Description
show	Show the time and date

Table 1-128 Options available with **lvm** command

Menu option	Description
add-disk	Add disk to the OS or data volume. You need to attach a disk before adding it. Table 1-129
list-free-disk	List the free disks
initialize-free-disk	Initialize the newly attached free disk
list-used-disk	List the disks used by the OS or data volume. Table 1-130

Table 1-128 Options available with **lvm** command (*continued*)

Menu option	Description
remove-disk	<p>Remove disk from the data volume. Remove disk operation involves migrating data from the existing disk to a new disk. You can remove a disk only after attaching a new disk with enough storage to migrate the data.</p> <p>The command first displays the list of disks being used and you need to select the disk that you want to remove. Then it displays the list of free disks where you want to migrate data and you need to select the disk. You can choose whether to initialize the new disk or not.</p> <p>It is recommended to suspend replication of all the configured Veritas Replication Sets before performing the remove disk operation.</p>
resize-logicalvolume	<p>Resize the OS or data volume for resized data disk</p> <p>Table 1-131</p>

Note: In case you initialize the newly-added disk during add-disk or remove-disk operation, the existing data on the new disk is deleted.

Table 1-129 Options available with **add-disk** command

Menu options	Description
data-volume	Add disk to the data volume
os-volume	Add disk to the OS volume

Table 1-130 Options available with **list-used-disk** command

Menu options	Description
data-volume	Lists disks used by the data volume
os-volume	Lists disks used by the OS volume

Table 1-131 Options with available with **resize-logicalvolume** command

Menu options	Description
data-volume	Resize the data volume
os-volume	Resize the OS volume

Table 1-132 Options available with **timezone** command

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

Table 1-133 Options available with **password-policies** command

Menu option	Description
set	Modify the administrator user password policies for the appliance.
show	Show the administrator user password-policies for the appliance.

Table 1-134 Options available with **password-policies set** command

Menu option	Description
max-age	Modify the maximum number of days before password change is required for administrator user.
min-age	Modify the minimum number of days before password change is required for administrator user.
min-length	Modify minimum password length for the administrator user.
warning-days	Modify number of days before a warning for administrator password expiry is given.

Table 1-135 Options available with **support** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
reset-support-password	Reset the support user password to the default installation password. This option may typically be used by Veritas support.
shell	Open the bash shell prompt for support user

Table 1-135 Options available with **support** command (*continued*)

Menu option	Description
loggather	If the appliance is configured as an IMS, then various options will be available for collecting the IMS logs. Table 1-136

Table 1-136 Options available with **loggather** command

Menu option	Description
basic	Gather logs of IMS without database See Table 1-137 on page 608.
full	Gather logs of IMS with database
cleanup	Clean up the loggater files of IMS
show	Lists all the loggater URLs ordered by date and time of IMS

Table 1-137 Options available with **basic** command

Menu option	Description
Number of days	Displays the basic logs from the days (1-99) mentioned

Table 1-138 Options available with **updates** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
Show-version	Show the appliance version
extra-drivers	Shows the list of the drivers and updates the extra drivers. Table 1-139
prepare-for-upgrade	Save the virtual appliance configuration in preparation for upgrade
rollback-update	Roll back the prepare for update operation

Table 1-139 Parameters needed for **extra-drivers** command

Menu option	Description
list	Show list of drivers and whether the updates available for that driver.
update	Update the drivers

Table 1-140 Options available with **utilities** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
clear	Clear the screen
unmount-cd-rom	Unmount the CD-ROM from the appliance
troubleshoot run-tool	Use the troubleshooting menu options Table 1-141
vmware-tools	Perform VMware Tools operations (install, uninstall, and show-version) Table 1-145
sftp-session	To start the SFTP server session Table 1-146
azure-waagent-service	Perform Azure waagent service operation. Applicable only in Azure environment Table 1-150
svc-delete-vdisk-timeout	To set or unset IBM SVC delete vdisk timeout value, which is the maximum amount of time that will be spent in retrying the delete vdisk operation during Cleanup Rehearsal. Table 1-151

Table 1-141 Options available with **troubleshoot run-tool** command

Menu option	Description
manage-nbu-primary-server-certificates	Manage the NetBackup certificates in Resiliency Platform Table 1-142
view-logs	View log files on any virtual appliance
check-port	Verify required open ports on Veritas Resiliency Platform VSA for communication with other appliance using admin password.

Table 1-142 Options available with **manage-nbu-primary-server-certificates** command

Menu option	Description
add	Add the certificate to the specified NetBackup primary server with token. Table 1-143
delete	Delete the certificate with specified fingerprint. Table 1-144
show	Show all the certificates.
help	Display the help text.

Table 1-143 Options available with **add** command

Menu option	Description
primary	Provide the NetBackup primary server hostname for reissuing the certificate
token	Provide the token value for the certificate registration

Table 1-144 Options available with **delete** command

Menu option	Description
fingerprint	Provide the fingerprint of specific certificate as an input to delete the certificate

Table 1-145 Options available with **vmware-tools** command

Menu option	Description
install	Install the VMware Tools mounted on CD-ROM of the appliance
show-version	Show the installed version of VMware Tools on the appliance
uninstall	Uninstall the VMware Tools from the appliance

Table 1-146 Options available with **sftp-session** command

Menu option	Description
start	To start the SFTP server session Table 1-147
show-details	View the current SFTP user and session details
stop	To stop the SFTP server session

Table 1-147 Options available with **start** command

Menu option	Description
get	View the file types that can be downloaded from the SFTP server Table 1-148
put	View the file types that can be uploaded on the SFTP server Table 1-149

Table 1-148 Options available with **get** command

Menu option	Description
logs	View the log files and directories from the SFTP server
heap-dump	Download the heap dump files of the service available on Resiliency Manager

Table 1-149 Options available with **put** command

Menu option	Description
patch	Upload the private patch on the SFTP server

Table 1-150 Options available with **azure-waagent-service** command

Menu option	Description
start	Start Azure waagent service
stop	Stop Azure waagent service
status	Show current status of Azure waagent service

Table 1-151 Options available with **svc-delete-vdisk-timeout** command

Menu option	Description
set	Sets SVC delete vdisk timeout value.
show	Shows SVC delete vdisk timeout value.
unset	Unsets SVC delete vdisk timeout value.

Klish menu options for Replication Gateway

Following are the options available for Replication Gateway using klish menu:

Table 1-152 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
hotfix	Manage hotfixes Table 1-153
manage	Manage appliance Table 1-154
monitor	Monitor appliance activities Table 1-160

Table 1-152 Options available in the **main** menu (*continued*)

Menu option	Description
network	Manage network configuration Table 1-163
settings	Manage appliance settings Table 1-172
support	To access logs Table 1-181
utilities	Run miscellaneous utilities Table 1-184

Table 1-153 Options available with **hotfix** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
apply-hotfix	Apply the specified hotfix
list-applied-hotfixes	List the applied hotfixes
list-available-hotfixes	List the available hotfixes
uninstall-hotfix	Uninstall the specified hotfix

Table 1-154 Options available with **manage** command

Menu option	Description
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
datamover	Manage Resiliency Platform Data Mover activities and objects Table 1-155
services	Manage the appliance services Table 1-157

Table 1-154 Options available with **manage** command (*continued*)

Menu option	Description
staging-storage	Perform Staging Storage Operations Table 1-158
cdp-storage	Manage CDP related storage Table 1-159

Table 1-155 Options available with **datamover** command

Menu option	Description
start	Start a Veritas Replication Set
abort	Stop a Veritas Replication Set
delete	Delete a Veritas Replication Set
resume	Resume a Veritas Replication Set
pause	Pause a Veritas Replication Set locally.
suspend-force	Suspend a Veritas Replication Set
clear-admin-wait	Clear the admin Wait status for the Veritas Replication Set
modify-quota-size	<p>Modify the size of the quota of a Veritas Replication Set. Modifying the quota affects the number of hosts that are protected with the gateway.</p> <p>You need to configure same quota size on all the peer gateways.</p> <p>Default quota size is 8000 MB. The minimum allowed quota size in direct mode is 2000 MB and in ObjectStore mode is 3500 MB. The maximum allowed quota size in direct and ObjectStore mode is 8000 MB.</p>
modify-updateset-size	<p>Modify the size of the UpdateSet. You need to configure same UpdateSet size on all the peer gateways.</p> <p>Default UpdateSet size is 500 MB. The minimum allowed UpdateSet size is 500 MB and the maximum allowed UpdateSet size is 2000 MB.</p>

Table 1-155 Options available with **datamover** command (*continued*)

Menu option	Description
modify-replication-frequency	<p>Modify the replication frequency.</p> <p>The default frequency is 120 sec.</p> <p>The minimum replication frequency allowed is 60 sec and maximum replication frequency allowed in 300 sec</p>
modify-tcpadvwin	<p>Modify the TCP socket's send and receive size between workload and the Replication Gateway. The default tcpadvwin value is 2 MB and maximum is 4MB. After the sizes are changed, replication of the consistency group has to be suspended and resumed to reflect this change.</p> <p>To suspend and resume the changes, use <code>suspend-force</code> and <code>resume</code> commands.</p> <p>Note: This command is not applicable for in-guest Windows workloads.</p>

Table 1-156 Options available with **fips** command

Menu option	Description
enable	Disable FIPS on the Replication Gateway appliance.
disable	Enable FIPS on the Replication Gateway appliance.
status	Show the current status of FIPs mode for the Replication Gateway appliance.

Table 1-157 Options available with **services** command

Menu option	Description
restart	<p>Restart Resiliency Platform services</p> <p>Two options available are:</p> <p><code>restart all</code> where, <i>all</i> means all the services.</p> <p><code>restart service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).</p>

Table 1-157 Options available with **services** command (*continued*)

Menu option	Description
start	<p>Start Resiliency Platform services</p> <p>Two options available are:</p> <p><code>start all</code> where, <i>all</i> means all the services.</p> <p><code>start service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).</p>
status	<p>Check the status of Resiliency Platform services</p> <p>Two options available are:</p> <p><code>status all</code> where, <i>all</i> means all the services.</p> <p><code>status service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).</p>
stop	<p>Stop Resiliency Platform services</p> <p>Two options available are:</p> <p><code>stop all</code> where, <i>all</i> means all the services.</p> <p><code>stop service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).</p>

Table 1-158 Options available with **staging-storage** command

Menu option	Description
add-disk	<p>Add disk to the staging-storage. You need to attach a disk before adding it.</p>
list-used-disk	<p>List the disks used in staging-storage.</p>

Table 1-158 Options available with **staging-storage** command (*continued*)

Menu option	Description
remove-disk	<p>Remove disk from the staging-storage. Remove disk operation involves migrating data from the existing disk to a new disk. You can remove a disk only after attaching a new disk with enough storage to migrate the data.</p> <p>The command first displays the list of disks being used and you need to select the disk that you want to remove. Then it displays the list of free disks where you want to migrate data and you need to select the disk. You can choose whether to initialize the new disk or not.</p> <p>It is recommended to suspend replication of all the configured Veritas Replication Sets before performing the remove disk operation.</p>
resize-logicalvolume	Resize the volume used in staging-storage.

Table 1-159 Options available with **cdp-storage** command

Menu option	Description
add-disk	To add a disk to the CDP storage
create-cdp-storage	To create a CDP storage
list-cdp-storage	To list all the CDP storage
list-used-disk	To list the disk used in the CDP storage
recover-cdp-storage	To recover the CDP storage
remove-cdp-storage	To remove the CDP storage
resize-logicalvolume	To resize the volume used in the CDP storage
recover-cdp-storage	To recover the CDP storage. This is also trigger data integrity test for the recovered CDP storage.
data-integrity-test	Perform data integrity test on data residing on respective CDP storage associated with Replication Set. Ensure that replication is not active for the opted replication set. One can perform Pause Table 1-155 before issuing this test.

Table 1-160 Options available with **monitor** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
top	Display the top process information
who	Display who is currently logged into the appliance
uptime	Display the uptime statistics for the appliance
FSuage	Display filesystem usage
datamover	Display VRP Datamover activities and objects Table 1-161

Table 1-161 Options available with **datamover** command

Menu option	Description
cdp-recovery-points	Display Veritas Replication Set details like recovery point ID, state, last received input /outputs, and time stamps
cdp-usage-stats	Displays details for each Veritas Replication Set like allocated space, free space, usage, etc.
usage-stats	Display the Resiliency Platform Data Mover usage stats
repl-sets	Display the details about Veritas Replication Sets including RPO, connection state, replication state. Display time required to sync the data and percentage of synced or replicated data. Table 1-161
update-sets	Display the list of current update sets which are in transit Table 1-161
ingress-data	Display the IO statistics for the data transfer from protected virtual or physical machine to Gateway (IOReceiver statistics) Table 1-161

Table 1-161 Options available with **datamover** command (*continued*)

Menu option	Description
network-data	Display the network related statistics for data transfer between production site Gateway and recovery site Gateway (Transceiver statistics) Table 1-161
disk-data	Display the IO statistics for the data write on recovery site disks (Applier statistics) Table 1-161
pair-status	Displays information and connectivity status of the local and the peer Gateway Table 1-161

Table 1-162 More information on the command

Datamover command	More information
repl-sets	<p>Name: Veritas Replication Set name for the protected virtual machine</p> <p>VRS-ID: Veritas Replication Set unique ID of a protected virtual machine</p> <p>Role: Role of the data center for the current Veritas Replication Set</p> <p>Data State: Replication data state for the current VRS-ID</p> <p>State: Replication State for the current VRS-ID</p> <p>Host Connection: Connection state of the protected virtual machine with the Replication Gateway</p> <p>Disks: Number of replicating disks for the protected virtual machine</p> <p>Lag (seconds): The time difference in seconds between a write operation occurs on the protected host on source side and the same gets replicated on the target side</p> <p>Admin Intervention: A flag notation if replication is broken. Check Admin wait state code</p> <p>Peer Gateway IP: IP address of the paired gateway</p>

Table 1-162 More information on the command (*continued*)

Datamover command	More information
update-sets	<p>Name: Veritas Replication Set name for the protected virtual machine</p> <p>VRS-ID: Veritas Replication Set unique ID of a protected virtual machine</p> <p>USID: Unique ID of the current update set. This is an increasing counter for each update set</p> <p>State: Current state of the Replication Update Set</p> <p>Size: Size of the data which is replicated in this update set</p> <p>Elapsed Time: Time of Update Set in the current state</p>
ingress-data	<p>VRS-ID: Veritas Replication Set Unique ID of a protected virtual machine</p> <p>USID: Unique ID of the current update set</p> <p>#Disk: Number of replicating disk or disks for protected virtual machine</p> <p>State: Shows different states of the virtual machine which is attached to the Replication Gateway. (INIT/DISCONNECTED/ACTIVE/SUSPENDED/ABORTED)</p> <p>Rate: IO rate of data from the protected host to Replication Gateway (Example: For the current update set, 907.0MB data has been received at the rate of 403.1Mb/s)</p> <p>Latency: This depicts the latency between the protected host to source Replication Gateway. Latency gives information about which component is slower</p> <p>Local deduplication: To know how much data is sent and the old data is cancelled (Amount of data written will be cancelled / total data size)</p>

Table 1-162 More information on the command (*continued*)

Datamover command	More information
network-data	<p>VRS-ID: Veritas Replication Set Unique ID of a protected virtual machine</p> <p>USID: Unique ID of the current update set</p> <p>Direction:</p> <p>Send: If update set is being sent from the Replication Gateway.</p> <p>Receive: If update set is being received by the Replication Gateway.</p> <p>Size: Size of the replicated data in this update set</p> <p>Rate: This is the rate at which the data is being written to the target disks by replication gateway</p> <p>Latency: Latency gives information about which component is slower</p> <p>Compression Ratio: Shows how much data is compressed while sending over WAN. The compression ratio is equal to actual data in update set/data sent over WAN</p>
disk-data	<p>VRS-ID: Veritas Replication Set Unique ID of a protected virtual machine</p> <p>USID: Unique ID of the current update set</p> <p>#Disk: Number of replicating disk or disks for protected virtual machine</p> <p>Size: Size of the replicated data in this update set</p> <p>Rate: This is the rate at which the data is being written to the target disks by replication gateway</p> <p>(Example: If update is 907MB, then rate at which it is written to the disk can be 403.1Mbps)</p> <p>Latency (Read, Write): Shows relative latency between staging area disk read and target disk write</p>

Table 1-162 More information on the command (*continued*)

Datamover command	More information
pair-status	<p>Pair Name: Name of the pair created between local and peer gateway</p> <p>Peer Gateway IP: IP address of the paired gateway</p> <p>Gateway Pair ID: Unique ID of the gateway which is paired</p> <p>Peer Gateway ID: Unique ID of the peer gateway</p> <p>Config State: Configuration state of the pair</p> <p>Connection State: Status of the connection between the local and the peer gateway</p>

Table 1-163 Options available with **network** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
dns	Show or set the DNS server or manage the options for resolv.conf file Table 1-164
ip	Show the IP address Table 1-166
route	View and manipulate the IP routing table Table 1-167
search-domain	Show or change the domain Table 1-168
traceroute	Trace packet routes to a particular host. You can also specify a port to trace the packet routes.
ssh-enabled-nic	Show or update SSH enabled NIC Table 1-169
nic-configuration	Show and configure the NIC Table 1-170

Table 1-163 Options available with **network** command (*continued*)

Menu option	Description
nic-for-UI	Show or update NICs configured to access product user interface Table 1-171

Table 1-164 Options available with **dns** command

Menu option	Description
options	Show, add, or remove options to the <code>/etc/resolv.conf</code> file. Refer to the documentation of <code>resolv.conf</code> for a list of available options and their purpose. Table 1-165
set	Configure Domain Name Server
show	Show the current Domain Name Server

Table 1-165 Options available with **options** command

Menu option	Description
add	Add a <code>resolv.conf</code> option
remove	Remove a <code>resolv.conf</code> option
show	Show options of <code>resolv.conf</code> file

Table 1-166 Options available with **IP** command

Menu option	Description
show	Show the current IP address

Table 1-167 Options available with **route** command

Menu option	Description
add	Set a default route or a route for a host or a subnet
delete	Delete the route entry from the routing table
show	Display your current routing table

Table 1-168 Options available with **search-domain** command

Menu option	Description
add	Add a search-domain
remove	Remove the search domain name
show	Show the search domain settings

Table 1-169 Options available with **ssh-enabled-nic** command

Menu option	Description
show	Show the NICs on which SSH is enabled. By default, SSH is enabled on all the NICs.
add	Add NIC to enable SSH on it
remove	Remove NIC to disable SSH on it

Table 1-170 Options available with **nic-configuration** command

Menu option	Description
show	Show details of NIC configuration like hostname, IPv4 or IPv6 address, prefix, gateway etc.
set	Configure the NICs which are not used while bootstrapping.

Table 1-171 Options available with **nic-for-UI** command

Menu option	Description
show	Show the NICs which are used to access product web user interface.
set	Set a NIC to access product web user interface from existing configured NICs.
remove	Remove one of the NICs used to access product web user interface.

Table 1-172 Options available with **settings** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session

Table 1-172 Options available with **settings** command (*continued*)

Menu option	Description
help	Display an overview of the CLI syntax
date	Display the current date and time for the appliance Table 1-173
lvm	Perform operations related to logical volume manager on the appliance Table 1-174
ntp	Perform operations related to NTP server
change-password	Change the admin user password for the appliance
poweroff	Shut down the appliance
fips	Enable, disable, or view the status of FIPS mode for a Replication Gateway Table 1-156
configure	Configure Resiliency Platform component or show the configured component
reboot	Restart the appliance
timezone	Show or change the timezone for the appliance Table 1-178
password-policies	Perform operation related to password policies of administrator user for the appliance Table 1-179

Table 1-173 Options available with **date** command

Menu option	Description
show	Show the time and date

Table 1-174 Options available with **lvm** command

Menu option	Description
add-disk	Add disk to the OS or data volume. You need to attach a disk before adding it. Table 1-175
list-free-disk	List the free disks
initialize-free-disk	Initialize the newly attached free disk
list-used-disk	List the disks used by the OS or data volume. Table 1-176
remove-disk	Remove disk from the data volume. Remove disk operation involves migrating data from the existing disk to a new disk. You can remove a disk only after attaching a new disk with enough storage to migrate the data. The command first displays the list of disks being used and you need to select the disk that you want to remove. Then it displays the list of free disks where you want to migrate data and you need to select the disk. You can choose whether to initialize the new disk or not. It is recommended to suspend replication of all the configured Veritas Replication Sets before performing the remove disk operation.
resize-logicalvolume	Resize the OS or data volume for resized data disk. Table 1-177

Table 1-175 Options available with **add-disk** command

Menu options	Description
data-volume	Add disk to the data volume
os-volume	Add disk to the OS volume

Table 1-176 Options available with **list-used-disk** command

Menu options	Description
data-volume	Lists disks used by the data volume
os-volume	Lists disks used by the OS volume

Table 1-177 Options available with **resize-logical-volume** command

Menu options	Description
data-volume	Resize the data volume
os-volume	Resize the OS volume

Note: In case you initialize the newly-added disk during add-disk or remove-disk operation, the existing data on the new disk is deleted.

Table 1-178 Options available with **timezone** command

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

Table 1-179 Options available with **password-policies** command

Menu option	Description
set	Modify the administrator user password policies for the appliance.
show	Show the administrator user password-policies for the appliance.

Table 1-180 Options available with **password-policies set** command

Menu option	Description
max-age	Modify the maximum number of days before password change is required for administrator user.
min-age	Modify the minimum number of days before password change is required for administrator user.
min-length	Modify minimum password length for the administrator user.
warning-days	Modify number of days before a warning for administrator password expiry is given.

Table 1-181 Options available with **support** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
reset-support-password	Reset the support user password to the default installation password. This option may typically be used by Veritas support.
shell	Open the bash shell prompt for support user
loggather	If the appliance has been configured as a Replication Gateway, then loggather full command will collect the logs of the Replication Gateway. Table 1-182

Table 1-182 Options available with **loggather** command

Menu option	Description
full	Gather logs of Replication Gateway with databaseSee Table 1-183 on page 628.
cleanup	Clean up the loggather files of Replication Gateway
show	Lists all the loggather URLs ordered by date and time of Replication Gateway

Table 1-183 Options available with **full** command

Menu option	Description
Number of days	Displays the full logs from the days (1-99) mentioned

Table 1-184 Options available with **utilities** command

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

Table 1-184 Options available with **utilities** command (*continued*)

Menu option	Description
clear	Clear the screen
unmount-cd-rom	Unmount the CD-ROM from the appliance
troubleshoot run-tool	Use the troubleshoot menu options Table 1-185
vmware-tools	Perform VMware Tools operations (install, uninstall, and show-version) Table 1-186
sftp-session	Use SFTP session for file transfer operation on the SFTP server Table 1-187
azure-waagent-service	Perform Azure waagent service operation. Applicable only in Azure environment
device-path-id	Show details of the attached disks to the Replication Gateway Table 1-188

Table 1-185 Options available with **troubleshoot run-tool** command

Menu option	Description
view-logs	View log files on any virtual appliance
check-port	Verify required open ports on Veritas Resiliency Platform VSA for communication with other appliance using admin password.

Table 1-186 Options available with **vmware-tools** command

Menu option	Description
install	Install the VMware Tools mounted on CD-ROM of the appliance
show-version	Show the installed version of VMware Tools on the appliance
uninstall	Uninstall the VMware Tools from the appliance

Table 1-187 Options available with **sftp-session** command

Menu option	Description
start	To start the SFTP server session Table 1-189
show-details	View the current SFTP user and session details
stop	To stop the SFTP server session

Table 1-188 Options available with **device-path-id** command

Menu option	Description
size	Show device path, disk ID, and size of the attached disks to the Replication Gateway

Table 1-189 Options available with start command

Menu option	Description
get	View the file types that can be downloaded from the SFTP server Table 1-190
put	View the file types that can be uploaded on the SFTP server Table 1-191

Table 1-190 Options available with **get** command

Menu option	Description
logs	Download the log files and directories from the SFTP server
heap-dump	Download the heap dump files of the service available on Resiliency Manager

Table 1-191 Options available with **put** command

Menu option	Description
patch	Upload the private patch on the SFTP server

Table 1-192 Options available with **azure-waagent-service** command

Menu option	Description
start	Start Azure waagent service
stop	Stop Azure waagent service
status	Show current status of Azure waagent service

About applying updates to Resiliency Platform

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

Updating the Resiliency Platform (Resiliency Manager and IMS) involves saving configuration to the data disk of the previous version appliance and then attaching the data disk to a new, freshly deployed virtual appliance.

For more details, refer the below topics:

About deploying the Resiliency Platform virtual appliances

Deployment workflows See [“Deployment workflows”](#) on page 370.

In Veritas Resiliency Platform, you can apply updates to the following:

- Veritas Resiliency Platform virtual appliance
- Veritas Resiliency Platform add-ons
- Host packages on the assets that are added to the Infrastructure Management Server (IMS) as a host
- Veritas Replication VIB

Upgrade path for Veritas Resiliency Platform 10.0

Veritas Resiliency Platform 3.6 is the minimum version supported for upgrade to version 10.0. If your installed version is older than minimum supported version, then you should first upgrade to your supported upgrade version.

Appliance upgrade sequence for Veritas Resiliency Platform 10.0

Appliances should be upgraded in following sequence:

1. All Resiliency Manager appliances
2. All IMS appliances
3. All Replication Gateway appliances

4. All managed hosts

Table 1-193 Upgrade path for Veritas Resiliency Platform 10.0

Current version	Can upgrade to
v3.6.0.0 with any hotfix applied	v10.0
v4.0.0.0 with any hotfix applied	v10.0

Considerations for applying update

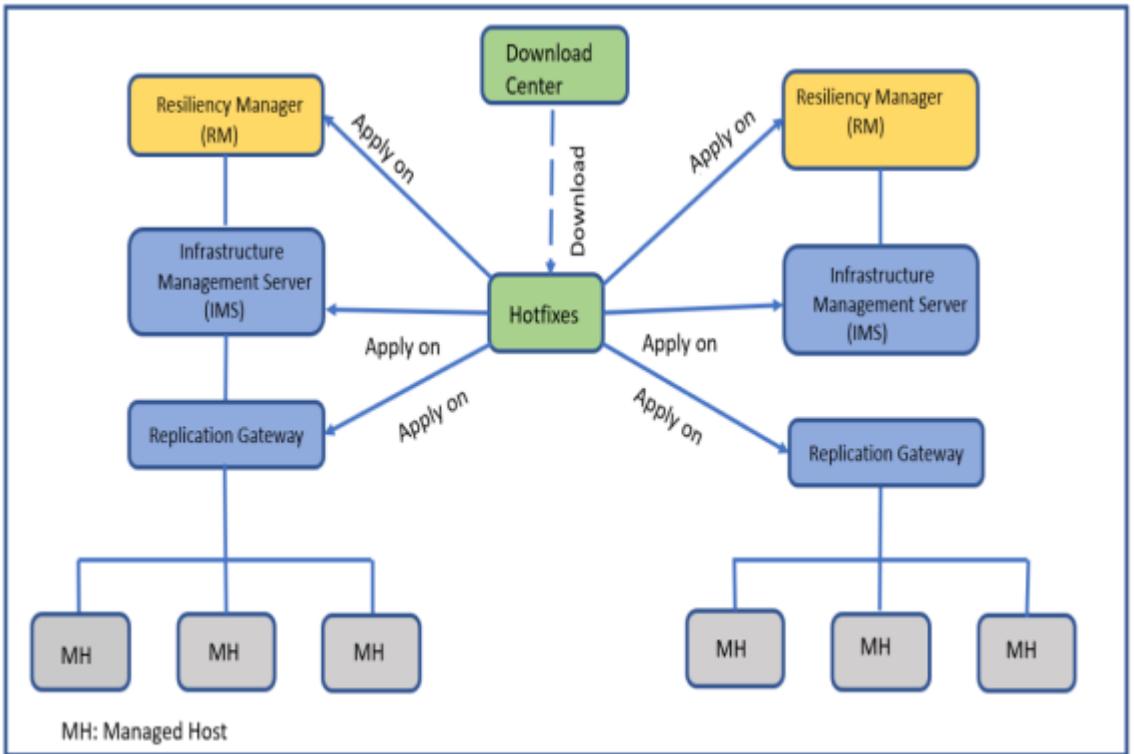
Following are some considerations for applying update to the Resiliency Platform components:

- The minimum version of the virtual appliances should be 3.6 or 4.0 to upgrade to Resiliency Platform 10.0.
- You must apply the updates on Resiliency Manager and IMS to take the complete advantage of the changes available in the updates. You need to replace the Replication Gateway to reflect the new upgrade changes in version 4.0.
- Ensure that no workflow or resiliency plans are running in the background.
- After applying updates, you can access the Resiliency Manager web user interface through the NIC that is used for communication with other Resiliency Manager.

You can change the NICs that are used to access product user interface. Use `network > nic-for-UI` Klish command to add additional existing configured NICs for UI access.

See [“Klish menu options for Resiliency Manager”](#) on page 586.

A pre-upgrade checklist is integrated with the `prepare-for-update` command to validate the virtual appliance. This checklist is applicable to Resiliency Manager only. For more details, refer See [“Pre-upgrade checklist”](#) on page 641.



The Resiliency Platform upgrade process is divided into 3 main parts. The upgrade process is applicable for the following platforms:

- Upgrading Resiliency Platform in AWS environment
- Upgrading Resiliency Platform in Azure environment
- Upgrading Resiliency Platform in VMware environment
- Upgrading Resiliency Platform in Hyper-V environment

The detail steps are mentioned according to the respective topics based on the environments. The high level overview of the process for applying updates to the virtual appliances of the Resiliency Platform in version 10.0 is mentioned below:

1. Download and install the required hotfix from Download Center to prepare the virtual appliances for upgrade.
2. Apply the hotfixes to the virtual appliances using KLIST commands to prepare the virtual appliances for upgrade.

See [“Klist menu options for Resiliency Manager”](#) on page 586.

See [“Klish menu options for IMS”](#) on page 598.

3. Detach the data disk of the existing virtual appliance and attach the data disk of the virtual appliance to the new 10.0 appliance.
4. Boot the virtual appliance to start the automatic bootstrap process based on the environment you want to upgrade.
5. In case of Replication Gateway upgrade, you need to replace the existing Replication Gateway with the newer version Replication Gateway.

Before applying updates, ensure that you have performed pre-upgrade tasks for specific scenarios:

See [“Pre-upgrade tasks”](#) on page 636.

The following table consists of the topics you need to refer to the steps mentioned above for applying updates in Veritas Resiliency Platform:

Table 1-194 Applying updates to Resiliency Platform

Step	Task	Steps to perform the task
1	<ol style="list-style-type: none"> a. Download the hotfix from Download Center. b. Apply the hotfix. 	<ol style="list-style-type: none"> a. See “Step 1: Downloading the Resiliency Platform update” on page 636. b.
2	<p>Prepare for update using Klish commands.</p> <p>Prepare the virtual appliances for update using Klish commands.</p>	<p>See “Step 2: Prepare for upgrade” on page 638.</p> <p>See “Klish menu options for Resiliency Manager” on page 586.</p> <p>See “Klish menu options for IMS” on page 598.</p>
3	<p>Detach / attach data disk.</p> <p>You need to perform following steps:</p> <ol style="list-style-type: none"> 1 Detach the data disk from the existing virtual appliance. 2 Create new virtual appliance and attach the existing data disk <p>Choose the environment from the list on which you wish to upgrade the Resiliency Platform.</p>	<p>See “Step 3: Upgrading the Resiliency Platform (Detach / attach the disk)” on page 642.</p>
4	<p>Start Automatic bootstrap process.</p> <p>After the bootstrap process is complete, the upgrade process should automatically start.</p>	<p>See “Step 4: Start the automatic bootstrap process” on page 659.</p>

Table 1-194 Applying updates to Resiliency Platform *(continued)*

Step	Task	Steps to perform the task
5	Replace the Replication Gateway appliance	See “Steps to replace the Replication Gateway appliance” on page 660. For modifying the encryption:
6	Apply update to the single or multiple Resiliency Managers in a domain. In case of multiple Resiliency Managers in the domain, the update needs to be applied on all the Resiliency Managers in synchronization. Note that if you apply update to the Resiliency Manager, then you must apply update to the IMS as well. If you do not update the IMS, the IMS stops data reporting to the Resiliency Manager.	See “Applying update on Resiliency Managers” on page 349.

Apart from upgrading the Resiliency Platform virtual appliances, you need to update following other components too:

Table 1-195 Applying updates to other components of Resiliency Platform

Task	Steps to perform the task
Apply update on the host packages. You must apply update to the Control Host after you apply update to the IMS.	
Apply update on the InfoScale environment In case of InfoScale environment, apply update to the add-on.	
Upgrading the Veritas Replication VIB In case of recovery of VMware virtual machines to on-premises data center using Resiliency Platform Data Mover, If you had configured a resiliency group before applying update to IMS, you need to apply update to the Data Mover bundle.	

Table 1-195 Applying updates to other components of Resiliency Platform
(continued)

Task	Steps to perform the task
Applying updates to the Veritas Data Gateway	

You also have an option of applying a private hotfix, if Veritas support provides you one.

Pre-upgrade tasks

Before upgrading, you may need to perform certain operations in certain scenarios.

Disconnected Resiliency Manager or IMS

If you are using Veritas Resiliency Platform and if any Resiliency Manager or IMS in your resiliency domain is in disconnected state, you must not upgrade. Contact Veritas support to fix the issue before you upgrade the Resiliency Platform components in your resiliency domain.

Rehearsal for recovery of VMware virtual machines to on-premises data center using Resiliency Platform Data Mover

If you have configured resiliency groups for recovery of VMware virtual machines to on-premises data center using Resiliency Platform Data Mover, ensure to perform cleanup rehearsal on the resiliency groups where you have performed rehearsal before applying update on Resiliency Manager.

Before upgrade, if you do not perform cleanup rehearsal on the resiliency groups where you have performed rehearsal, the cleanup rehearsal operation completes successfully after upgrade but other disaster recovery operations such as migrate or recover fails.

See [“About applying updates to Resiliency Platform”](#) on page 631.

Step 1: Downloading the Resiliency Platform update

This topic is a part of the main topic that explains the end-to-end process of applying an update to Veritas Resiliency Platform components. To understand the sequence in which the update needs to be applied to various Veritas Resiliency Platform components, you must see:

Table 1-196 Hotfix details which can applied on

Resiliency Platform version	Hotfix number	Hotfix description	Download location
For version 3.6	3.6.0.20	<p>This is a mandatory hotfix to be applied on Resiliency Platform to upgrade to v10.0. It needs to be installed on all Resiliency Manager, IMS, and Replication Gateways before executing Klish command:</p> <pre>updates > prepare-for-update</pre> <p>While prepare-for-update is in progress, this hotfix ensures that all the pending processes are complete. It monitors the pending processes for at least 30 minutes. If there are any pending processes in incomplete state prepare-for-update will exit after 30 minutes.</p> <p>For more details refer Pre-upgrade checklist</p>	Hotfix location on Download Center
For version 4.0	4.0.0.6	<p>This is a mandatory hotfix to be applied on Resiliency Platform to upgrade to v10.0. It needs to be installed on all the Resiliency Manager, IMS, and Replication Gateways before executing Klish command <code>updates>prepare-for-update</code>.</p> <p>While prepare-for-update is in progress, this hotfix ensures that all the pending processes are complete. It monitors the pending processes for at least 30 minutes. If there are any pending processes in incomplete state prepare-for-update will exit after 30 minutes.</p> <p>For more details refer Pre-upgrade checklist</p>	Hotfix location on Download Center

To download the update for v3.6, follow below mentioned steps:

- 1 To apply the hotfix, download the .tar.gz files from the above mentioned locations.
- 2 Ensure that you do not unzip the .tar.gz files before uploading it to the virtual appliance. Copy the downloaded files to the required location of the appliance and apply the hotfix.

To download the update for v4.0, follow below mentioned steps:

- 1 To apply the hotfix, download the .hotfix files from the above mentioned locations.
- 2 Ensure that you do not change .hotfix files before uploading it to the virtual appliance. Copy the downloaded files to the required location of the appliance and apply the hotfix.

Next step is to **Prepare for upgrade**, where it is required to prepare the virtual appliance for upgrade.

More Information

Step 2: See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 3: See [“Step 3: Upgrading the Resiliency Platform \(Detach / attach the disk\)”](#) on page 642.

Step 4: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Step 2: Prepare for upgrade

The upgrade process is divided into 3 major parts for all the virtual appliances in which the `Prepare for upgrade` is the first step. You need to perform below steps to prepare the virtual appliances for upgrade using the klish commands:

Prerequisite:

Make sure you have downloaded the appropriate hotfix before upgrading Resiliency Platform to the latest version.

Table 1-197 Hotfix details which can applied on Resiliency Platform

Resiliency Platform version	Hotfix number	Hotfix description	Download location
For version 3.6	3.6.0.20	<p>This is a mandatory hotfix to be applied on Resiliency Platform to upgrade to v10.0. It needs to be installed on all Resiliency Manager, IMS, and Replication Gateways before executing Klish command:</p> <pre>updates> prepare-for-update.</pre> <p>While prepare-for-update is in progress, this hotfix ensures that all the pending processes are complete. It monitors the pending processes for at least 30 minutes. If there are any pending processes in incomplete state prepare-for-update will exit after 30 minutes.</p> <p>For more details refer Pre-upgrade checklist</p>	Hotfix location on Download Center
For version 4.0	4.0.0.6	<p>This is a mandatory hotfix to be applied on Resiliency Platform to upgrade to v10.0. It needs to be installed on all the Resiliency Manager, IMS, and Replication Gateways before executing Klish command</p> <pre>updates> prepare-for-update.</pre> <p>While prepare-for-update is in progress, this hotfix ensures that all the pending processes are complete. It monitors the pending processes for at least 30 minutes. If there are any pending processes in incomplete state prepare-for-update will exit after 30 minutes.</p> <p>For more details refer Pre-upgrade checklist</p>	Hotfix location on Download Center

Steps to perform upgrade

- 1 Login to KLISH menu of the RM appliance and apply the desired hotfix using below command.

```
hotfix > apply-hotfix
```

For example: `hotfix > apply-hotfix <hotfix_number>`

Note: If your data center consists of multiple Resiliency Managers, you have to apply the hotfix on all the RMs.

- 2 On KLISH menu of the RM appliance, execute the below command to prepare the RM for upgrade.

```
updates > prepare-for-update
```

A pre-upgrade checklist is integrated with the `prepare-for-update` command to validate the virtual appliance. This checklist is applicable to Resiliency Manager only. For more details, refer See [“Pre-upgrade checklist”](#) on page 641.

Note: If your data center consists of multiple Resiliency Managers, you have to execute this command on any one RM appliance. This command will prepare all the other RMs for the upgrade. It will also take the backup of all the configurations and power off all the RMs.

- 3 Wait until the RM appliance is powered off automatically. On the RM appliance console, a message appears as “The system is going down for power-off now.

To upgrade the IMS appliance to the latest version, repeat the steps 1-3 on all the available IMS appliances.

The next step is **Upgrading the Resiliency Platform (Attach / detach the disk)** where it is required to detach the data disk of the previous version virtual appliances and attach it to the current version virtual appliance.

See [“Step 3: Upgrading the Resiliency Platform \(Detach / attach the disk\)”](#) on page 642.

Since the upgrade procedure is different for the environments supported by the Resiliency Platform, the steps are different for each environment.

More Information

- Step 1:** See [“Step 1: Downloading the Resiliency Platform update”](#) on page 636.

Step 3: See “[Step 3: Upgrading the Resiliency Platform \(Detach / attach the disk\)](#)” on page 642.

Step 4: See “[Step 4: Start the automatic bootstrap process](#)” on page 659.

Apply updates [About applying updates to Resiliency Platform](#)

Klish menu for Resiliency Manager [Klish menu options for Resiliency Manager](#)

Klish menu for IMS [Klish menu options for IMS](#)

Pre-upgrade checklist

To make upgrade experience smoother a pre-update validation checklist is added as part of prepare-for-update to check and find any discrepancy before starting the upgrade process. This checklist is applicable to Resiliency Manager only.

Following checks are done in pre-upgrade checklist:

1. **Verify pre-upgrade configuration:**

■ **Duplicate objects:**

If the configuration consists of duplicate objects, the post upgrade operations might fail. It checks the database and fetches if any duplicate objects are found. If duplicate objects are found, a warning is displayed.

If the duplicate objects persist, the post-upgrade operations fail. It checks the database and fetches if any duplicate objects are found. If duplicate objects are found, this check will fail and you need to clean the duplicate object with help of Veritas Support and retry.

■ **Stale objects:** These objects might be present in the environment due to any of the following reasons:

- If any stale Consistency Groups present.

- If there is any failed rehearsal operation present.

- If you have performed “force cleanup” which might have left residue. Stale objects block the replace Replication Gateway as part of upgrade activity. To reduce this downtime, these are validated as part of pre-upgrade checks. These stale objects do not block the upgrade process, but it is recommended to clean these duplicate objects with help of Veritas Support. This can be done either before upgrade process or just after upgrading the Resiliency Manager.

2. **Check for on-going activities:** While upgrade is in progress, number of on-going activities is displayed. These activities might get terminated if you confirm to proceed with the upgrade.

3. **Check for on-going reports:** While upgrade is in progress, number of on-going reports is displayed. These activities might get terminated if you confirm to proceed with the upgrade.
4. **Check pending process:** While upgrade is in progress, multiple number of process are pending for execution. You need to make sure these processes are completed or the upgrade operation may exit. While upgrade is in progress, make sure all the processes are complete. If multiple number of process are pending for execution, prepare-for-update will exit after 30 minutes.

See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 3: Upgrading the Resiliency Platform (Detach / attach the disk)

The upgrade process is divided into 3 major parts for all the virtual appliances in which the detach / attach data disk is the second step. In the below mentioned environments, you can detach the data disk of the previous version virtual appliance and attach it to the current / latest version virtual appliance.

- [Upgrading Resiliency Platform in VMware environment](#)
- [Upgrading Resiliency Platform in Hyper-V environment](#)
- [Upgrading Resiliency Platform in AWS environment](#)
- [Upgrading Resiliency Platform in Azure environment](#)

Next step is to **Start the automatic bootstrap process** for the virtual appliances.

More Information:

Step 1: See [“Step 1: Downloading the Resiliency Platform update”](#) on page 636.

Step 2: See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 4: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates: See [“About applying updates to Resiliency Platform”](#) on page 631.

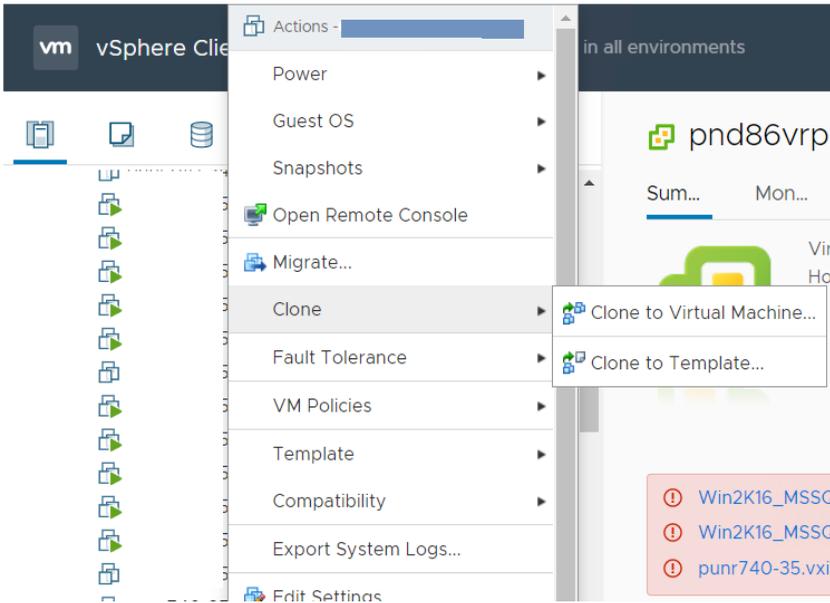
Upgrading Resiliency Platform in VMware environment

The upgrade process is divided into 3 major parts for all the virtual appliances in which the **Attach / Detach data disk** is the second step. You need to perform the below steps to detach the data disk from the previous version virtual appliance and attach it to the current version virtual appliance in the vCenter server:

To attach / detach the data disk

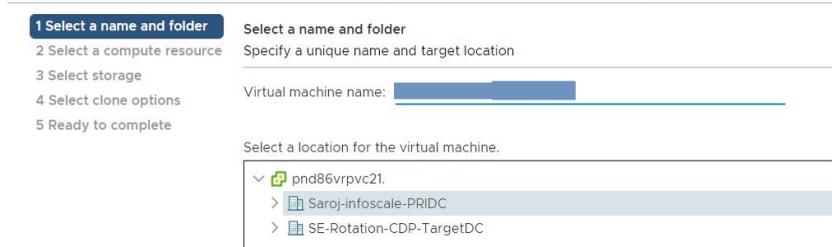
- 1** Ensure that the prepare for upgrade is performed on the virtual appliance. Refer [Step 2: Prepare for upgrade](#).
- 2** Login to vSphere client using admin user credentials.
- 3** Verify that the virtual appliance on which you have applied the updates are shut down.

- 4 To clone the virtual appliance of the previous version perform the below steps:
 - a. Right click on the virtual appliance and select the **Clone > Clone to Virtual Machine** option.



- b. Provide user friendly Virtual machine name to the cloned virtual machine.

- Clone Existing Virtual Machi...



- c. Click **Next**.
 - d. Select the compute resource and click **Next**.
 - e. Select the storage and click **Next**. Complete the clone virtual machine process.

punr [redacted] - Clone Existing Virtual Machi...

1 Select a name and folder
 2 Select a compute resource
 3 Select storage
 4 Select clone options
 5 Ready to complete

Select storage
 Select the storage for the configuration and disk files

Select virtual disk format: Same format as source
 VM Storage Policy: Keep existing VM storage policies

Configure per disk

Name	Capacity	Provisioned	Free	Type
PrimDS1	4.36 TB	5.15 TB	2.14 TB	VM <input type="button" value="v"/>

- 5 Perform the cloning steps on the another virtual appliance.
- 6 Deploy the current version virtual appliance. Do not complete the bootstrap process. To deploy the virtual appliance perform the steps 1-10 from the topic, refer [Deploying the virtual appliance through VMware vSphere Client](#)

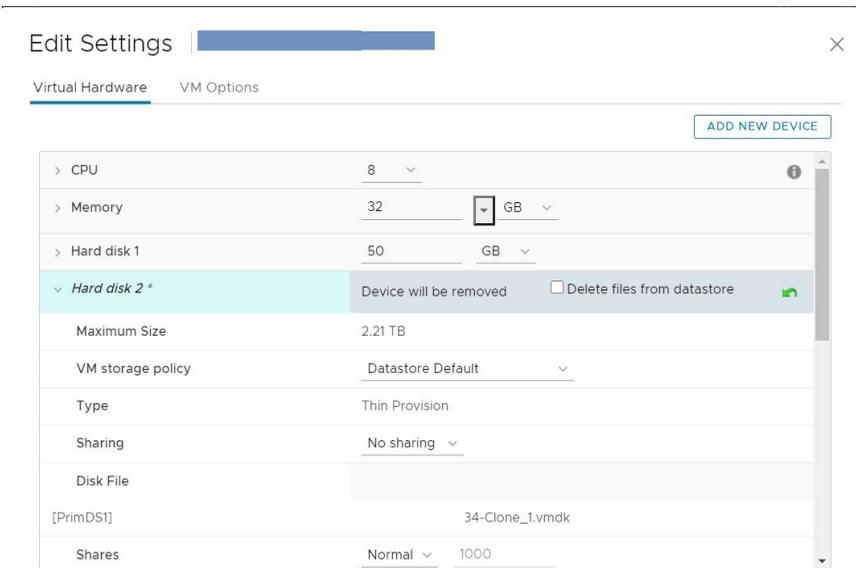
7 The next step is to detach the data disk of the previous version from the appliance. Perform the following steps while you are logged into the vCenter server:

- a. Right click on the appliance and select **Edit settings**.
- b. Select the data disks and choose 'X' to detach the disk.

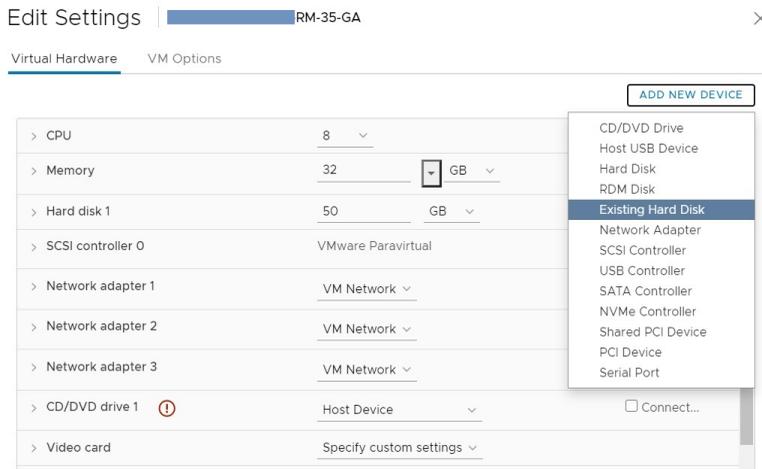
Note: Do not select the ‘Delete files from datastore’ check box “. This option will delete the selected disk.

c. From the **Disk File** section, note the path of the data disks which is detached in step 7b.

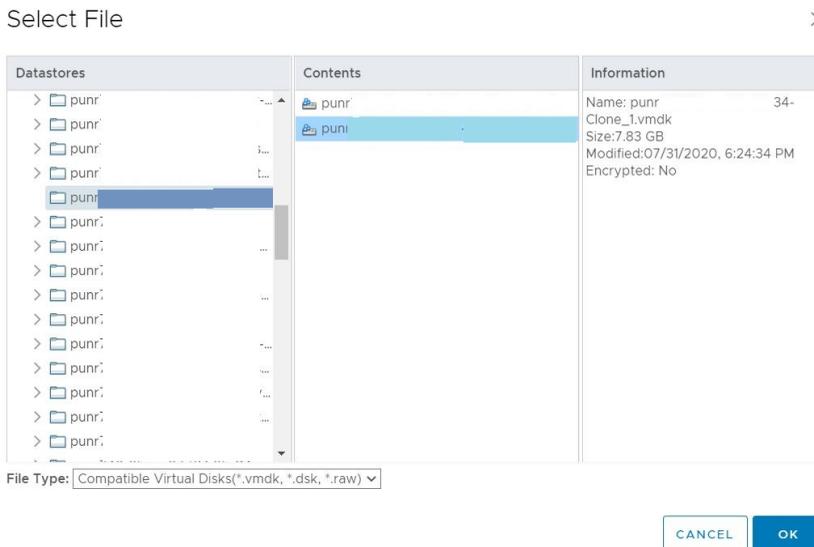
Perform these steps on the other virtual appliances.



- 8 Next step is to attach the data disk to the current version (4.0) virtual appliance which is deployed in step 6 above. Perform the following steps:
 - a. Right click on the virtual appliance and select **Edit settings**.
 - b. Click on the **Add New Device** button and select **Existing Hard Disk** option.



- c. Select the datastore and then select the data disk which is noted in step 7c. Click **OK**.



Perform the steps 8a – 8c on the another virtual appliance to attach the data disk.

- 9 The last step is to power on the virtual appliance. Perform the below steps:
Right click on the appliance and click the **Power On** option. You need to power on the other appliances too.

Note: Attaching the data disk from one virtual appliance to another is supported only during the upgrade process, i.e. after preparing the virtual appliance for upgrade step ([Step 2: Prepare for upgrade](#)). If the data disk of an appliance is changed during normal functioning, it may impact the DR operations.

More Information

Previous Step: See [“Step 2: Prepare for upgrade”](#) on page 638.

Next Step: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

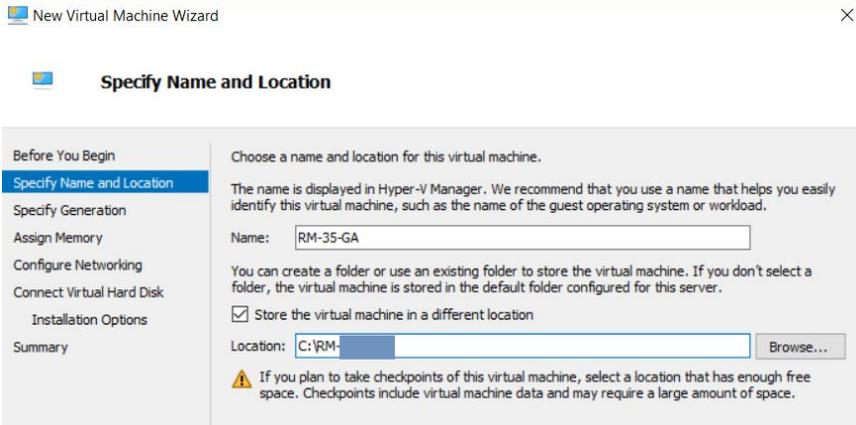
Upgrading Resiliency Platform in Hyper-V environment

The upgrade process is divided into 3 major parts for all the virtual appliances in which the **Create new virtual appliance and attach the data disk** is the second step. You need to perform below steps to create a new appliance and attach the previous version virtual appliance data disk to the current version virtual appliance in the Hyper-V server:

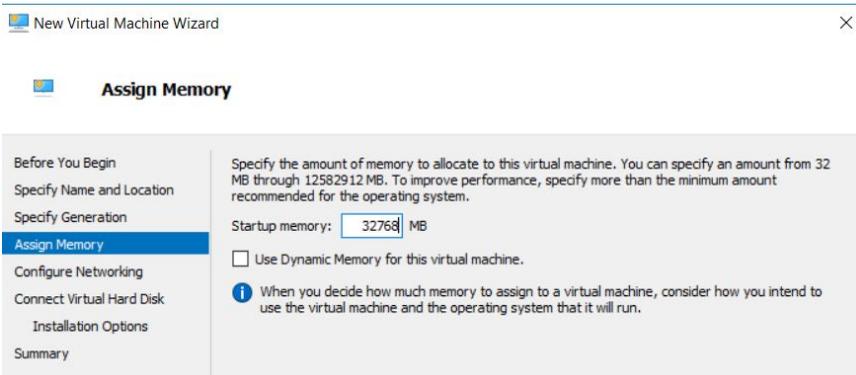
Create new virtual appliance and attach the existing data disk

- 1 Ensure that the prepare for upgrade is performed on the virtual appliance.
Refer [Step 2: Prepare for upgrade](#)
- 2 Login to Hyper-V Manager client using admin user credentials.

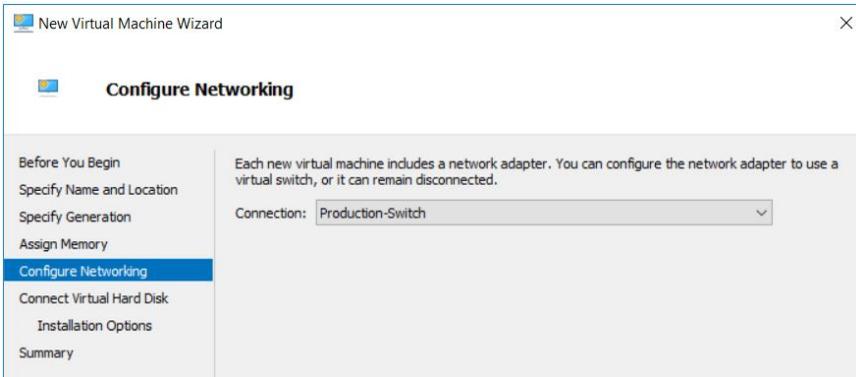
- 3 To deploy the current version virtual appliance, perform the following steps:
 - a. Right-click on the virtual machine and select **New > Virtual Machine**. This opens up the **New Virtual Machine Wizard**.
 - b. Provide the name and location for the virtual machine on the **New Virtual Machine Wizard** and click **Next**.



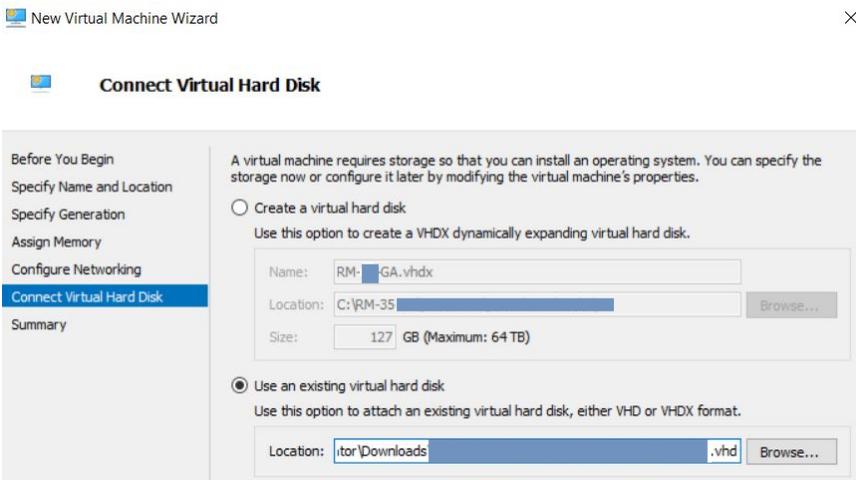
- c. Select the **Generation 1** and click **Next**.
 - d. Provide **Startup memory** in digits based on the system requirement of the Resiliency Platform, refer and click **Next**.



- e. Provide **Configure Networking** as **Production-Switch** and click **Next**.

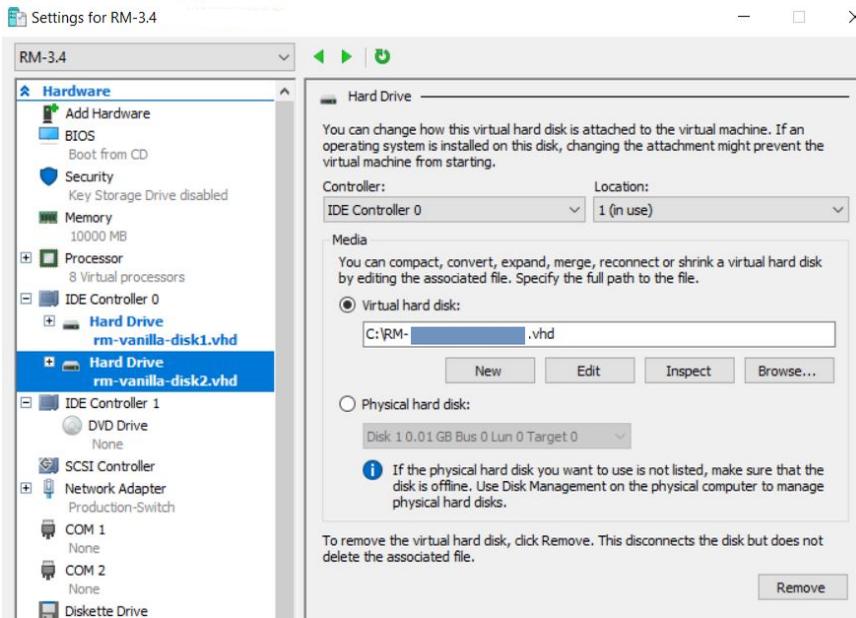


f. In **Connect Virtual Hard Disk** wizard, select **Use an existing virtual hard disk** option and browse to the location where you have saved the version 4.0 hard disk and select Hard Disk 1. Click **Finish**.

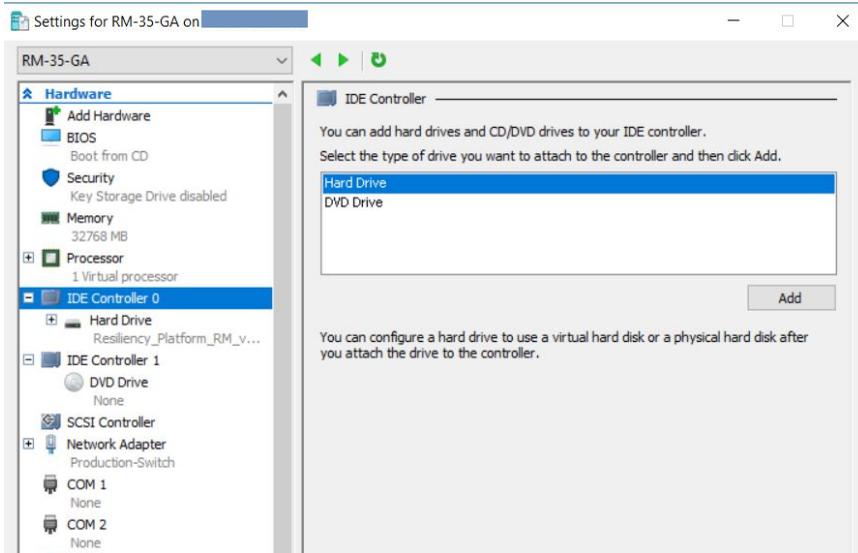


4 Right click on the previous version (3.5 or 3.6) virtual appliance and select **Settings >> IDE Controller >> Hard Disk**. On the **Virtual hard disk** section, select the **Browse** to locate the existing data disk.

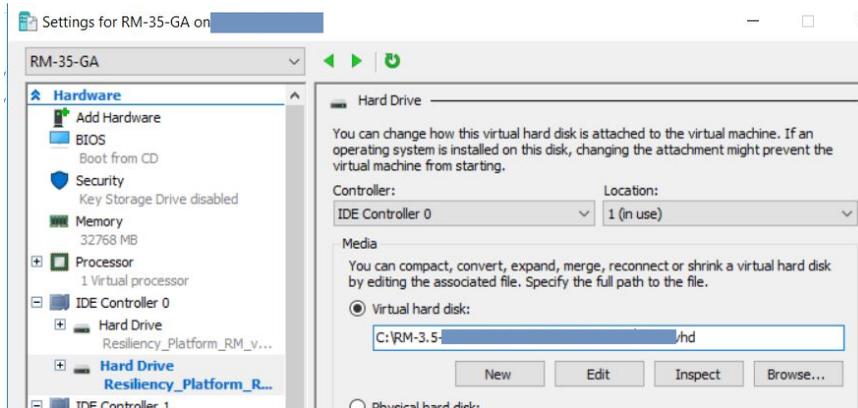
You need to copy the previous (3.5 or 3.6) data disk to the new 4.0 virtual appliance and note the new path of the data disk,



- 6 a. Go to the **Settings** of the 4.0 virtual appliance, select IDE Controller 0, select **Hard Drive** and click **Add**.



- b. Browse the copied data disk location noted in Step 4.



- 7 Click **Apply** after it is enabled and click **OK**.
- 8 Start the virtual appliance by right clicking on the appliance and click **Start** option.

Note: Attaching the data disk from one virtual appliance to another is supported only during the upgrade process, i.e. after preparing the virtual appliance for upgrade step ([Step 2: Prepare for upgrade](#)). If the data disk of an appliance is changed during normal functioning, it may impact the DR operations.

More Information

Previous Step: See [“Step 2: Prepare for upgrade”](#) on page 638.

Next Step: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

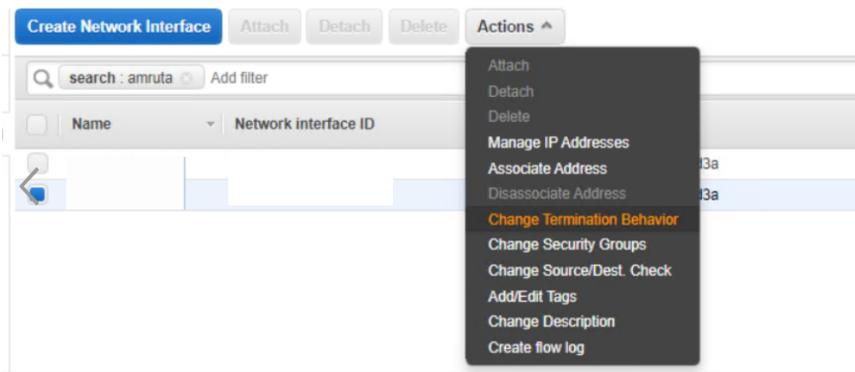
Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Upgrading Resiliency Platform in AWS environment

The upgrade process is divided into 3 major parts for all the virtual appliances in which the **Attach / Detach data disk** is the second step. You need to perform below steps to detach the data disk from previous version virtual appliance and attach it to the version current virtual appliance.

Detach / attach data disk

- 1 Login to AWS portal along with the admin user credentials.
- 2 Navigate to EC2 service and click **Network Interfaces** on the left pane, search for the desired IP address. Provide the name to the NIC and make note of the private IP addresses of the appliance.
- 3 To preserve the NICs of the corresponding IP addresses, navigate to:
 - a. **Actions** drop down.
 - b. Select the **Change Termination Behavior** option.
 - c. Uncheck the **Delete on termination** checkbox and click **Save**.

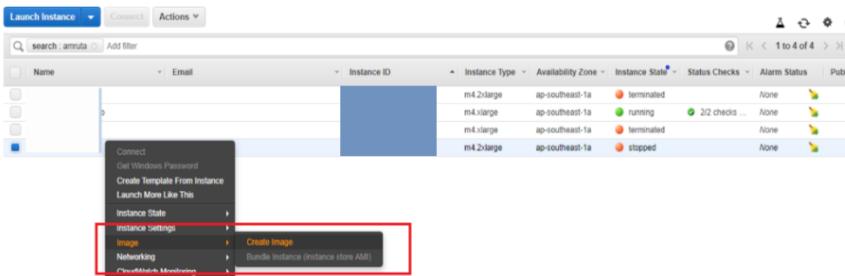


- 4 Make sure you add port 7000 for Resiliency Manager in existing Security Group.
- 5 Navigate to the **Instances** on the left pane. To create an AMI image of the instance i.e backup of the existing version image, navigate to **Actions > Image > Create Image**.

Provide the following details on the **Create Image** wizard and click on the **Create Image**.

Enter Image name:

Enter Image description:



We need to perform this step to roll back the version AMI image if something goes wrong during the upgrade process. For rollback operation, refer the topic See [“Rollback steps in AWS environment”](#) on page 662. .Wait for the Image creation to complete before going to the next page.

- 6 Detach all data disks from the previous version appliance and recommended to tag them with a proper name.
- 7 Delete the older (previous) version virtual appliance to reuse the private IP addresses while configuring the new appliances. Perform the following steps:
 - a. Right-click on older version virtual appliance, select the **Instance State > Terminate**.
 - b. Click on **Yes, Terminate** on the **Terminate Instances** wizard.

- 8 Go to the **AWS Marketplace** and locate **Veritas Resiliency Platform** product. Choose the option of **Upgrade** to upgrade the appropriate virtual appliance.

Parameters
Parameters are defined in your template and allow you to input custom values when you create or update a stack.

Resiliency Manager EC2 Instance Configuration

Instance Name
Enter a name for the new Resiliency Manager EC2 instance
VRP_RM2

EC2 Instance Type
Select a valid AWS instance type from the list for the new Resiliency Manager instance, or leave as is to automatically select the recommended instance type based on the appliance and region
--- Automatically select best fit instance type ---

Key Pair for SSH access
Select an existing EC2 key pair from the list that will be used to enable SSH access to the new Resiliency Manager EC2 instance

Data Volumes from Previous Version
Select ALL data volumes that were attached to the older Resiliency Manager instance. If any volume is missed it would cause issues in the upgrade process.

Network Configuration

VPC ID
Select the VPC in which the new Resiliency Manager EC2 instance will be deployed

Network interface to attach as eth0
Enter the Network interface id of the eth0 network interface, if applicable

Network interface to attach as eth1
Enter the Network interface id of the eth1 network interface, if applicable

Cancel Previous Next

- 9 Provide the required inputs:
 - Instance Name:** Provide appropriate name for the new version instance.
 - EC2 Instance Type:** Select the appropriate instance type to be used.
 - Key Pair for SSH access:** Select the appropriate key pair for the EC2 instance.
 - Data Volumes from Previous Version:** Select all data volumes that were detached from previous version virtual appliance from the drop down list.Under **Network Configuration** provide the below details:
 - VPC ID:** Provide the VPC id of the previous version virtual appliance.
 - Network interface to attach as eth0:** Enter NIC ids that were noted from previous virtual appliance.
 - Network interface to attach as eth1:** Enter NIC ids that were noted from previous virtual appliance.
- 10 The new version virtual appliance is deployed. Using SSH, continue to complete the upgrade steps.

Note: Attaching the data disk from one virtual appliance to another is supported only during the upgrade process, i.e. after preparing the virtual appliance for upgrade step ([Step 2: Prepare for upgrade](#)). If the data disk of an appliance is changed during normal functioning, it may impact the DR operations.

More Information

Previous Step: See [“Step 2: Prepare for upgrade”](#) on page 638.

Next Step: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Upgrading Resiliency Platform in Azure environment

The upgrade process is divided into 3 major parts for all the virtual appliances in which the **Attach / Detach data disk** is the second step. You need to perform below steps to detach the data disk from the previous version virtual appliance and attach it to the current version virtual appliance.

Detach / attach disk

- 1 Login to Azure console and navigate to the appliance. Click on the **Capture** option.



- 2 Fill the required details and create an image. This image can be used for the rollback in case the upgrade process fails for any reason.
- 3 Before deleting previous version appliances, note the hostname of the appliance. Use the same hostname while deploying the upgrade offerings.
- 4 Delete the previous version virtual appliance. Ensure that the NICs and volumes of the deleted appliances are not deleted along with the virtual appliance.

- Go to the **Azure Marketplace** and deploy the offering **Veritas Resiliency Platform** by providing following values:

Basic settings for Azure deployment

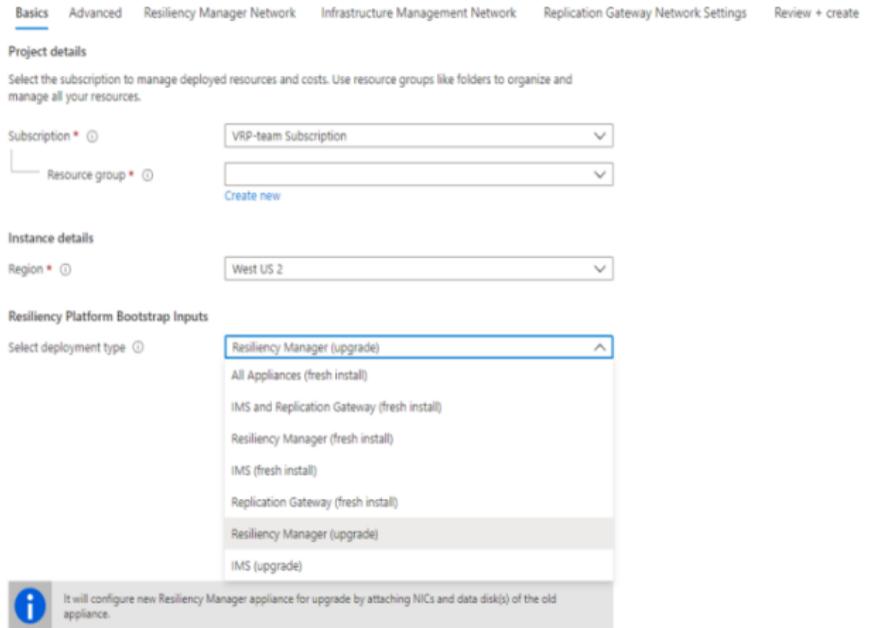


Table: Basic settings for Azure deployment

Input field	Description
Select Deployment Type	Select the appropriate deployment type from the given options. Select 'Resiliency Manager (upgrade)' or 'IMS (upgrade)' for upgrading Resiliency Manager and IMS respectively.
Password for admin user	Set the password for admin user. The admin user and password is later used for configuring the appliances.
Confirm password	Provide same password for confirmation.
Subscription	Select the subscription of Azure account, to be used for deploying the virtual appliances.
Resource group	Specify the name of an existing resource group that contains disks and NICs for the existing appliance, which is being upgraded.

Advance settings for Azure deployment



Table: Advance settings for Azure deployment

Input field	Description
Instance Name	Name of newer version appliance.
Hostname	Hostname for the newer appliance. This is used if custom DNS is not configured appropriately. Make sure that the hostname is same as the older appliance which is noted before deleting the previous version appliance.
Instance Size	Size of newer version appliance.
Existing Data Disk Name(s)	Provide comma separated list of data disk(s) in same sequence as the LUN number.

Network settings for Azure deployment

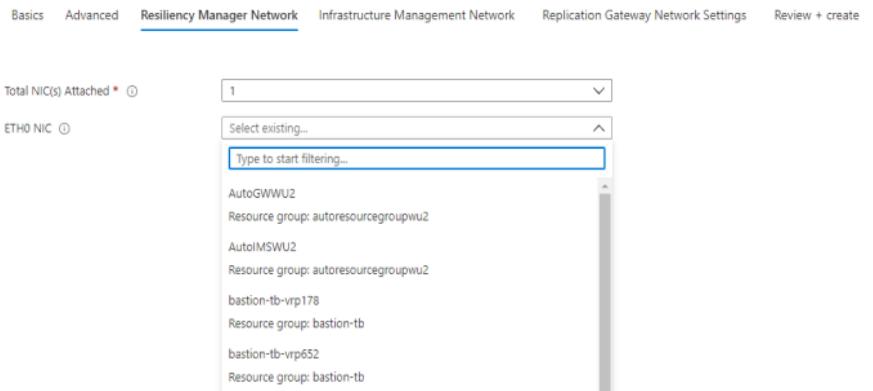


Table: Network settings for Azure deployment

Input field	Description
Total NIC(s) Attached	Select number of NICs attached to older version appliance.
ETHX NIC	Select the appropriate existing NIC attached to the older appliance.

- 6 Click **OK** and review the summary displayed in the **Summary** section.
- 7 Click **OK** to view and accept the terms and conditions.
- 8 Click **Create** to create the instances:
 - This step creates newer version virtual appliance with existing data disk(s) and NICs. To complete the upgrade login to the console with 'admin' username and provided password and continue with further configuration. Make sure you add port 7000 for Resiliency Manager in existing Security Group.

Note: Attaching the data disk from one virtual appliance to another is supported only during the upgrade process, i.e. after preparing the virtual appliance for upgrade step ([Step 2: Prepare for upgrade](#)). If the data disk of an appliance is changed during normal functioning, it may impact the DR operations.

More Information

Previous Step: See [“Step 2: Prepare for upgrade”](#) on page 638.

Next Step: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Rollback steps in Hyper-V environment

The rollback operation returns the image or any action return to the previous state. Hence, you may also require rollback steps in case the upgrade operation fails.

Steps to rollback in Hyper-V environment

- 1 Delete the newer version appliance which you have deployed in step 3 in the topic See [“Upgrading Resiliency Platform in Hyper-V environment”](#) on page 648.
- 2 Power on the older version appliance.
- 3 Run `update > prepare-for-update` command on the older appliance which reverts back the configuration as before.

More Information

Apply Updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Step 4: Start the automatic bootstrap process

The upgrade process is divided into 3 major parts for all the virtual appliances in which the `Prepare for upgrade` is the first step. Detaching the data disk and attaching it to the new virtual appliance is the next step. The last step is to start automatic bootstrap process. Below are the steps:

Start the automatic bootstrap process

- 1 While you are in the vSphere client, launch the web console of the RM appliance.
- 2 On the login prompt, provide the default password. Change the password and repeat the same step on the other RM appliances.
- 3 If the appliance detects a data disk from a previous version, you can see a screen showing the network configuration details. Confirm that these details correspond to the correct appliance and proceed to start the bootstrap process.
- 4 Complete the bootstrap process. The upgrade process should start automatically.
- 5 To upgrade the Resiliency Manager refer to [Applying update on Resiliency Managers](#).

In case of multiple Resiliency Managers in the domain, the update needs to be applied on all the Resiliency Managers in synchronization.

Perform steps 1-5 on all the IMSs as well.

Note: If you apply update to the Resiliency Manager, then you must apply update to the IMS as well. If you do not update the IMS, the IMS stops reporting data to the Resiliency Manager.

Once the RM and IMS virtual appliances are upgraded to latest version, the Replication Gateway needs to be replaced with the new Replication Gateway appliance. Refer See [“Steps to replace the Replication Gateway appliance”](#) on page 660.

For more information on upgrading Resiliency Managers, refer the section called “Applying update on Resiliency Managers”

More Information:

Step 1: See [“Step 1: Downloading the Resiliency Platform update”](#) on page 636.

Step 2: See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 3: See “[Step 3: Upgrading the Resiliency Platform \(Detach / attach the disk\)](#)” on page 642.

Apply updates: See “[About applying updates to Resiliency Platform](#)” on page 631.

Applying update on Resiliency Managers

If you have multiple Resiliency Managers in your resiliency domain then you need to apply the update on all the Resiliency Managers in synchronization.

For the purpose of upgrade, one of the Resiliency Managers is designated as import node and the other one as non-import node:

- The first Resiliency Manager configured in the resiliency domain is called the import node. If the first Resiliency Manager is no longer in the resiliency domain, then the Resiliency Manager configured next in the domain is designated as the import node.
- The other Resiliency Manager in the domain is called the non-import node.

If you do not remember the sequence in which the Resiliency Managers are configured in the domain, you can start applying the update on any one of the Resiliency Managers and the process will guide you about the import node and non-import node.

To apply update on a single Resiliency Manager in the domain

- 1 Login to Resiliency Manager virtual appliance using admin credentials.
- 2 Click **Yes** when asked to verify the network configuration settings.
- 3 The bootstrap process for configuring the Resiliency Manager starts.
- 4 After the configuration of the RM is complete, the upgrade step to the new version starts along with verifying status for some of the below mentioned points:
 - Updating configuration file with NIC details.
 - Updating the product version.
 - Syncing the data between the Resiliency Managers.
- 5 After starting all the Resiliency Platform services, the upgrade process is complete.
- 6 A message appears **Upgrade completed from <old_version> to <new_version> successfully.**

To apply update on Resiliency Managers if you have multiple Resiliency Managers in the domain

- 1 Start the process of applying update on the import node. Ensure the health checkup for multiple RMs scenario. Make sure all the Resiliency Manager services are stopped. You are prompted to switch to one of the non-import nodes.
- 2 Start the process of applying update on the non-import nodes. You can apply update on the non-import nodes simultaneously. Wait until upgrade on all the non-import nodes reach to a stage where you are prompted to switch to the import node.
- 3 On import node, continue the process of applying update. Once **Performing schema and data sync up** stage gets completed on the import node, you are prompted to switch to the non-import nodes.
- 4 On the non-import nodes, continue the process of applying update in a sequential manner. The **Performing schema and data sync up** stage gets completed on all the non-import nodes.
- 5 Once again, continue the process of applying update on all the non-import nodes in a sequential manner. The **Applying new security** stage gets completed on all the non-import nodes. Once this stage gets completed on all the non-import nodes, you are prompted to switch to the import node once again and continue the process of applying update.
- 6 Now, the **Applying new security** stage gets completed on the import node and all the Resiliency Manager services are started.
- 7 Switch to the non-import nodes once again and enter **yes** to start the Resiliency Manager services on the non-import nodes.

Note: Do not perform the **Performing schema and data sync up** and **Applying new security** stages on the non-import nodes simultaneously.

Updating LDAP / AD configuration

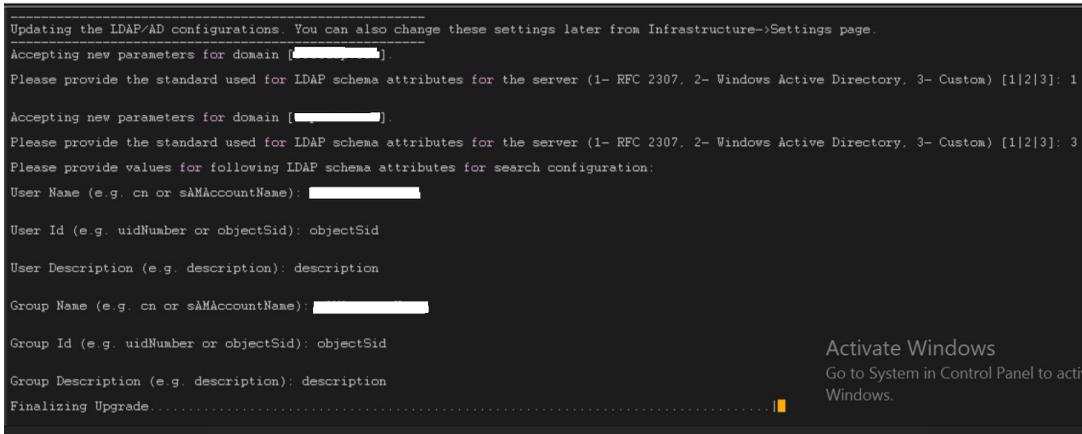
This procedure is applicable only when you have already configured LDAP / AD configuration in the previous versions. Based on the configuration of the LDAP / AD server, you need to select the appropriate standard. Only in case of custom standard, you need to provide values for some of the attributes. This steps are executed after starting the Resiliency Manager services.

You can also change these settings from **Infrastructure settings > Settings** page. For doing this, you may need the admin credentials to login which is extra step. Also this may overcome the purpose of initial settings done before the upgrade.

Hence it is expected to provide right values during Resiliency Manager upgrade so that LDAP / AD server user login works smoothly.

Figure 1-10 Updating LDAP / AD configuration

```
Updating the LDAP/AD configurations. You can also change these settings later from Infrastructure->Settings page.
Accepting new parameters for domain [REDACTED].
Please provide the standard used for LDAP schema attributes for the server (1- RFC 2307, 2- Windows Active Directory, 3- Custom) [1|2|3]: 1
Accepting new parameters for domain [REDACTED].
Please provide the standard used for LDAP schema attributes for the server (1- RFC 2307, 2- Windows Active Directory, 3- Custom) [1|2|3]: 3
Please provide values for following LDAP schema attributes for search configuration:
User Name (e.g. cn or sAMAccountName): [REDACTED]
User Id (e.g. uidNumber or objectSid): objectSid
User Description (e.g. description): description
Group Name (e.g. cn or sAMAccountName): [REDACTED]
Group Id (e.g. uidNumber or objectSid): objectSid
Group Description (e.g. description): description
Finalizing Upgrade..... [REDACTED]
```



See [“About applying updates to Resiliency Platform”](#) on page 631.

Steps to replace the Replication Gateway appliance

Once the RM and IMS virtual appliances are upgraded to the latest version, the Replication Gateway needs to be replaced with the new Replication Gateway appliance. There are two ways to replace the Replication gateway after the RM and IMS are upgraded.

1. Replace the Replication Gateway with same IP address and hostname
2. Replace the Replication Gateway with different IP address and hostname

For more information related to the prerequisites of replace Replication Gateway operation, refer [Replacing a Replication Gateway from a gateway pair](#)

Below are the steps to replace the existing gateway with a new Replication Gateway:

Replace the Replication Gateway with same IP address and hostname

Note: Make sure you do not replace the Replication Gateway and its peer gateway from the Replication Gateway pair simultaneously.

- 1 Login to vSphere client and switch off the existing version of Replication Gateway from source data center by right clicking on the Replication Gateway and select **Power Off** option.
- 2 Next step is to deploy the new Replication Gateway using the same IP address and hostname. You have to select in which environment you are deploying the Replication gateway and accordingly refer to the following topics:
 - [Deploying the virtual appliance through VMware vSphere Client](#)
 - [Deploying the virtual appliance through Hyper-V Manager](#)
 - [Deploying the virtual appliances in AWS through AWS Marketplace](#)
 - [Deploying the virtual appliances in AWS using OVA files](#)
 - [Deploying the virtual appliances in Azure using PowerShell script](#)
 - [Deploying the virtual appliances in Azure through Azure Marketplace](#)
 - [Deploying the virtual appliances in Google Cloud Platform using OVA files](#)
 - [Deploying the virtual appliances in Google Cloud Platform \(GCP\) through GCP Marketplace](#)
- 3 Login into RM and navigate to **Settings > Infrastructure > Data Mover** card mentioned path to check the status of the Replication Gateway is **Faulty**.
- 4 Add the new Replication Gateway which was deployed in step 2. Perform following steps:
 - a. Click on **+ Add Replication Gateway**.
 - b. Provide the IP address and password. Click **Submit**. The workflow is initiated.
- 5 The status of the Replication Gateway is **Healthy** after the workflow is complete.
- 6 Right click on the Replication Gateway and select **Replace Gateway** option.
- 7 Click **OK** on the **Replace Gateway** panel. The Replace Gateway workflow is initiated.
- 8 After the workflow is complete, navigate to the **Settings > Infrastructure > Data Mover** card.
- 9 Verify that the old Replication Gateway is replaced with the new Replication Gateway.
- 10 Navigate to **Assets > Replication Appliance**. The gateway pair status is in **Connected** state.

Perform the same steps to replace the Replication Gateway on the target data center.

For Klish options, See ["Klish menu options for Replication Gateway"](#) on page 612.

Replace the Replication Gateway with different IP address and hostname

- 1 Login to vSphere client and switch off the existing version of Replication Gateway from source data center by right clicking on the Replication Gateway and select **Power Off** option.
- 2 Next step is to deploy the new Replication Gateway using the different IP address and hostname. You have to select in which environment you are deploying the Replication gateway and accordingly refer to the following topics:
[Deploying the virtual appliance through VMware vSphere Client](#)
- 3 To replace the Replication Gateway, repeat the steps 3-10 from the above procedure **Replace the Replication Gateway with same IP address and hostname**.

While upgrading the Replication Gateways from older versions to latest version, it is recommended to replace the Replication Gateway. As part of this process, Replication Gateway with the older version needs to be powered off and should not be powered-on. Hence one of the recommendation is to perform log-gather of the Replication Gateway on older gateway appliance before replacing it.

More Information

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Replacing a Replication Gateway from a gateway pair

Using the Resiliency Platform console you can replace an existing Replication Gateway from a pair irrespective of whether the gateway is healthy or faulted.

To replace a gateway, you must go to the data center having the gateway and perform the replace gateway operation. In this case the peer gateway must be healthy.

If you want to replace both the gateways, then after replacing the first gateway wait till the operation is successfully completed before you replace the other gateway. If both the Replication Gateways are faulted, you cannot perform the replace Replication Gateway operation. You need to remove them and create a new pair.

Note: 1. Make sure you do not replace the Replication Gateway and its peer gateway from the Replication Gateway pair simultaneously. This may lead to replace Replication Gateway failure and you will have to delete the Replication Gateway, the gateway pair and its associated resiliency groups.

2. While replacing the Replication Gateway on the VMware environment, DRS automation level for the existing Replication Gateway and new Replication Gateway should be set to manual or should be disabled.

3. If you have replaced a Replication Gateway which is associated with a resiliency group which is protecting a physical workload, it is mandatory to perform resync operation once the replace gateway operation is successful.

Also, you can replace a gateway with another gateway only if it is configured to use the same mode of replication. For example, you cannot replace a gateway that is configured to use Direct mode of replication with the one that is configured to use Object Storage mode of replication.

When you replace a gateway, all the Veritas Replication Sets are moved to a new gateway and the Replication Gateway pairs are reconfigured. If the existing Replication Gateway is configured as a PXE boot server or as a DHCP server, then the new gateway is reconfigured as a PXE Boot server or a DHCP server.

Note that after replacing a gateway, irrespective of the replaced gateway being faulted or healthy, the replication resumes from the last successfully replicated update set.

Before you replace a replication gateway, if any of the resiliency groups that are associated with the gateway are in offline state, because the migrate or take over operation has failed, then you need to start the resiliency group before replacing the gateway. Ensure that you select the **Refresh storage, network, compute and customizations** check box in the start resiliency group wizard.

If the staging disk attached to the Replication Gateway is corrupted or removed, you need to shutdown the gateway and then perform the replace gateway operation.

After replacing a gateway, the **Server Gateway** attribute is inherited from the previous pair. If the Replication Gateway being replaced itself is the Server Gateway, then the attribute is updated to point to the new gateway. Ensure that the networking constraints between the two gateways are maintained with the new gateway.

Replace Replication Gateway operation is not supported in the following scenarios:

- When recovering VMware virtual machines from an on-premises data center to vCloud Director.
- When recovering virtual machines from vCloud Director to vCloud Director.

From version 3.5, you can configure the resiliency group for Continuous Data Protection (CDP) by selecting the **Enable Continuous Data Protection** checkbox while selecting the appropriate gateway pair.

To replace a Replication Gateway from a gateway pair

- 1 Navigate to **Settings** (menu bar) > **Infrastructure** > **Details View**
You can also access this page from the **Quick Actions** menu.
- 2 Click **Data Mover** > **Resiliency Platform Data Mover** tab.
- 3 In the **Replication Gateways** table, right-click the Replication Gateway which you want to replace, and select **Replace Gateway**.
- 4 From the list of gateways, select the replacement gateway, and click **Next** to initiate the replace operation.

The replace Replication Gateway operation may fail if there are any resiliency groups in configuration failed state, that are associated with the gateway. You then need to delete the Veritas Replication Sets, then delete the resiliency groups, and then retry the replace gateway operation.

While replacing a Replication Gateway having CDP storage attached, make sure that the datastore on which the CDP storage resides is accessible to new Replication Gateway.

Rollback steps in AWS environment

The rollback operation returns the image or any action return to the previous state. Hence, you may require rollback steps in case the upgrade operation fails.

Steps for rollback image in AWS environment

- 1 Login into the AWS portal.
- 2 Navigate to the virtual appliance of that version on which the error had occurred. Click on the **Delete** option, so that the NICs are free from the new virtual appliance. As performed in the step 6 of topic See [“Upgrading Resiliency Platform in AWS environment”](#) on page 653., preserve the NICs instead of new ones.
- 3 Launch the new instance using the previous version AMI which was noted in step 4 from the topic See [“Upgrading Resiliency Platform in AWS environment”](#) on page 653.. While launching the instance similar to Step 3 in topic See [“Upgrading Resiliency Platform in AWS environment”](#) on page 653., select the preserved NICs instead of creating the new NICs. Also select or create an appropriate IAM role with required permissions as mentioned in See [“Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway”](#) on page 389..

- 4 After the instance is launched, connect it using the admin user and previous version appliance password / SSH key credentials.
- 5 Login into Klish menu of the virtual appliance and execute the below command:

```
updates > rollback-update
```

The previous version virtual appliance is up and online.

More Information

Step 1: See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 2: See [“Upgrading Resiliency Platform in AWS environment”](#) on page 653.

Step 3: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Rollback steps in Azure environment

The rollback operation returns the image or any action to the previous state in case the upgrade fails.

Steps for rollback image in Azure environment

- 1 Download the `Veritas_Resiliency_Platform_Azure_Upgrade_Scripts_<version>.zip` from the Veritas Download Center.
- 2 Extract the bundle to obtain the `rollbackTemplate.json` and saved it locally.
- 3 Login to Azure portal.
- 4 Browse to the section **Deploy a custom template** option from Azure portal.
- 5 Select the option **Build your own template in the editor**.

- 6 Select the **Load file** option and load the `rollbackTemplate.json` file.
- 7 While clicking the **Save** button, it display the below template:

The screenshot shows the 'Custom deployment' configuration page in Google Cloud Platform. The page title is 'Custom deployment' with a subtitle 'Deploy from a custom template'. Below the title, there is a brief instruction: 'Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.' The form contains several sections:

- Subscription:** A dropdown menu with a red highlight over the selected option.
- Resource group:** A dropdown menu with a 'Create new' link below it.
- Instance details:**
 - Region:** A dropdown menu set to 'East US'.
 - Admin Password:** A text input field.
 - Appliance Name:** A text input field.
 - Appliance Hostname:** A text input field.
 - Appliance Role:** A dropdown menu set to 'Resiliency Manager'.
 - Old Appliance Version:** A dropdown menu set to '3.6.0.0'.
 - Appliance Deployment Type:** A dropdown menu set to 'marketplace'.
 - Appliance Vm Size:** A text input field.
 - Image Name:** A text input field.
 - Eth0Nic Name:** A text input field.
 - Eth1Nic Name:** A text input field.
 - Primary Nic:** A dropdown menu set to 'eth0'.

Below are the details of the attributes:

Table 1-198 Custom deployment template attributes

Attributes	Description
Subscription	Select the Subscription Name/ID where the original version Resiliency Platform virtual appliances are deployed.

Table 1-198 Custom deployment template attributes (*continued*)

Attributes	Description
Resource Group	This should be the same "Resource group" name that was used by the original Resiliency Platform virtual appliances.
Region	This is purely on customer discretion where they want to use.
Admin Password	This password is used by the Resiliency Platform virtual appliance once the rollback deployment is complete. It can be the same password used by the admin on the original virtual appliance (preferred) or provide a new password.
Appliance Name	This would be the display name for the Resiliency Platform virtual appliance
Appliance Hostname	<p>This would be the FQDN or the exact full name of the Resiliency Platform virtual appliance when it was of the original version.</p> <p>Eg: If the original Resiliency Platform Resiliency Manager name was azure-rm-pr.abc.local with display name as Azure-RM, then in that case, the "Appliance name" field will have the value as "Azure-RM" and the "Appliance hostname" will have the value as "azure-rm-pr.abc.local".</p>
Appliance Role	This can either be "Resiliency Manager" or "Infrastructure Management Server" depending on which type of appliance is planned for rollback.
Old Appliance Version	This would be current_version, for example if the older/original version was 4.0.
Appliance Deployment Type	This would be "Marketplace" if the original Resiliency Platform virtual appliance was deployed via Azure Marketplace itself.
Appliance virtual machine size	Ensure to select Standard_D8s_v3 for RM and Standard_F8s for IMS.

Table 1-198 Custom deployment template attributes (*continued*)

Attributes	Description
Image Name	This will be the image name that was created as a part of the “Prepare for upgrade” process. Refer to Apply updates >> Upgrading Resiliency Platform in Azure environment section in Product documentation.
Eth0Nic Name and Eth1Nic Name	Name of the NIC interfaces used by the older/original virtual appliance.
Primary NIC	This would be the “NIC” primarily used for all communications from the virtual appliance. By default, this is eth0 for all Resiliency Platform virtual appliance. (However, ensure to double check and confirm eth0 is the Primary NIC that was used by the older/original virtual appliance.)

8. Once the deployment process is completed, power on the virtual machine and login with the admin user with the password given in the above template.

9. Once successfully logged in, run the command `updates > rollback-update`.

Note: If the rollback is done in a multiple Resiliency Manager setup with above steps, the database rebuild might take some time and until then the “Database service” will be in STOPPED state. It may take up to 30 minutes for database rebuild / repair operation depending upon the size of the database. Once the process is completed, verify the database services using `manage > services status` as ALL.

10. Finally, login to the Resiliency Manager console to confirm all resiliency group details are reflecting as expected.

Step 1: See “[Step 2: Prepare for upgrade](#)” on page 638.

Step 2: See “[Upgrading Resiliency Platform in Azure environment](#)” on page 656.

Step 3: See “[Step 4: Start the automatic bootstrap process](#)” on page 659.

Apply updates See “[About applying updates to Resiliency Platform](#)” on page 631.

Rollback steps in VMware environment

The rollback operation returns the image or any action return to the previous state. Hence, you may require to rollback the steps in case the upgrade operation fails.

Steps to rollback in VMware environment

- 1 Delete the newer version appliance which you have deployed in step 6 in the topic [Upgrading Resiliency Platform in VMware environment](#)
- 2 Power on the older version appliance.
- 3 Run `update > prepare-for-update` command on the older appliance which reverts back the configuration as before

More Information:

Apply Updates: [About applying updates to Resiliency Platform](#)

Rollback steps in Hyper-V environment

The rollback operation returns the image or any action return to the previous state. Hence, you may also require rollback steps in case the upgrade operation fails.

Steps to rollback in Hyper-V environment

- 1 Delete the newer version appliance which you have deployed in step 3 in the topic See [“Upgrading Resiliency Platform in Hyper-V environment”](#) on page 648.
- 2 Power on the older version appliance.
- 3 Run `update > prepare-for-update` command on the older appliance which reverts back the configuration as before.

More Information

Apply Updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Virtual appliance security features

Veritas Resiliency Platform is shipped and deployed in the form of virtual appliances. Following are the security features of Veritas Resiliency Platform virtual appliances:

See [“Operating system security”](#) on page 668.

See [“Management Security”](#) on page 668.

See [“Network security”](#) on page 668.

See [“Access control security”](#) on page 669.

See [“Physical security”](#) on page 668.

Physical security

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

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.

The Control + Alt + Delete key combination has been disabled to avoid any accidental reboot of the virtual appliance. Exec-shield is enabled to protect the virtual appliance from stack, heap, and integer overflows.

Management Security

Only two users are available on the appliance: admin user and support user. These two user accounts are used to access the appliance based on the requirement.

Only admin login is available for the appliance. The password policy of admin login is modified to prompt the user to change the password on the first login.

See [“Password policies for Resiliency Platform virtual appliance”](#) on page 451.

If the admin user password is lost, you need to contact Veritas support for resetting the admin user password.

On successful completion of the product bootstrap, admin user can only access a limited menu of commands through klish. Besides admin user, support user is also supported in the appliance but remote login of support user is disabled. To access the support user, one need to login as an admin and go through **klish**. An option `support > shell` is provided in the **klish** menu to switch the user to support and access the bash shell of support. After selecting this option, the support user is given superuser privileges. Using this option is not recommended and it should be used only with the assistance of technical support.

Timeout of the bash shells of all users is set to 900 seconds.

Network security

The TCP timestamp responses are disabled in Resiliency Platform virtual appliance. Another network security feature of the appliance is that during the product bootstrap process, only those ports that are used by the product for communication and data transfer, are opened through the firewall and all the other communications are blocked.

Uncommon network protocols such as DCCP, SCTP, RDC, TIPC have been disabled so that any process cannot load them dynamically.

Access control security

Resiliency Platform virtual appliance implements certain access control measures. The umask is set to 0700 across the appliance. The access permissions of some of the files such as home folder of root, the log directory etc. is restricted. All the security and the authorization messages are logged into the appliance.

Recovering Hyper-V virtual machines to Google Cloud Platform

This chapter includes the following topics:

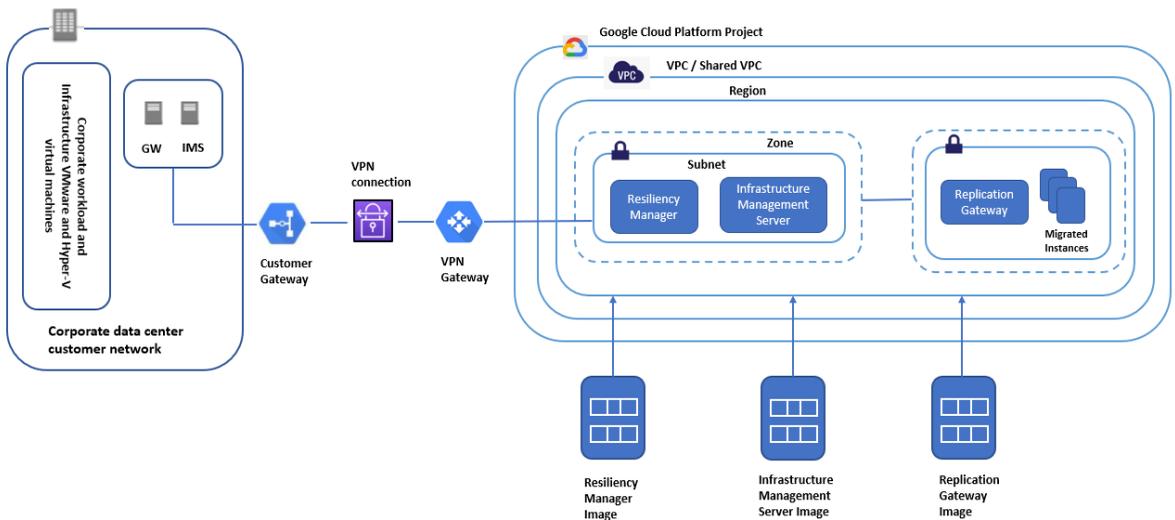
- [About recovering virtual machines to Google Cloud Platform](#)
- [Plan your environment](#)
- [Deploy and configure the virtual appliances](#)
- [Set up the resiliency domain](#)
- [Add asset infrastructure](#)
- [Infrastructure Pairing](#)
- [Create resiliency groups](#)
- [Advanced features](#)
- [Perform remote recovery operations](#)
- [Monitor assets](#)
- [Miscellaneous references](#)

About recovering virtual machines to Google Cloud Platform

Using Veritas Resiliency Platform, you can configure and protect your VMware virtual machines for recovery to Google Cloud Platform using the Resiliency Platform Data Mover. Following are the capabilities supported for this use case:

- Operations like Rehearsal, Recover, Migrate (and Migrate back), bulk-recovery are enabled when target data center is Google Cloud Platform.
- Network customization feature is added where you can enable or disable IP customization.
- Support for Shared VPC, which is a unique feature in Google Cloud Platform.
- Support for Regional disks, in addition to Zonal disks.
- Support for Customer Managed Encryption Keys.

Figure 2-1 Overview of deployment Infrastructure for recovery to Google Cloud Platform



Plan your environment

This topic explains about the way you are planning to configure the Resiliency Platform components in your environment. Refer to the **Overview and Planning**

Guide on SORT to know about the product, its components, features, and capabilities. Refer to the **Release Notes** on SORT for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.

How to find Resiliency Platform documents on SORT:

- 1 Navigate to [SORT home](#).
- 2 In the **Knowledge Base** tab, click on **Documents** link.
- 3 Select **Resiliency Platform** from the Product list. You can view all the links of the released versions.
- 4 Click on the **Unix or Windows** link to view all the guides.

Deploy and configure the virtual appliances

Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the Google Cloud Platform cloud data center as well as in the premises data center.

Refer to the following topics:

Deploy and configure virtual appliances in Google Cloud Platform :

To deploy the virtual appliances in the Google Cloud Platform region, refer the below sequence:

- See [“Downloading the Veritas Resiliency Platform virtual appliances”](#) on page 366.
- See [“About deploying the Resiliency Platform virtual appliances”](#) on page 368.
- See [“Deploying the virtual appliances in Google Cloud Platform \(GCP\) through GCP Marketplace”](#) on page 421.
- See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Deploy and configure virtual appliances in premise data center:

To deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center,

- See [“Deploying the virtual appliance through VMware vSphere Client”](#) on page 419. Deploying the virtual appliance through VMware vSphere Client

To configure the virtual appliances as Veritas Resiliency Platform components:

- See [“Prerequisites for configuring Resiliency Platform components”](#) on page 438.

- See “[About configuring the Resiliency Platform components](#)” on page 437.

Downloading the Veritas Resiliency Platform virtual appliances

You can download a licensed copy of the Veritas Resiliency Platform virtual appliances from [MyVeritas portal](#).

You can download the files for deploying virtual appliances. The virtual appliances are available in two formats: in the form of Open Virtualization Archive (OVA) files, or in the form of zip files. The .zip files contain the Virtual Hard disk (VHD) image file using which you can deploy the virtual appliances.

Table 2-1 Filenames for Veritas Resiliency Platform 10.0

Component	Filenames
Resiliency Manager	Veritas_Resiliency_Platform_RM_VMware_Virtual_Appliance_10.0.0.0_IE.ova
IMS	Veritas_Resiliency_Platform_IMS_VMware_Virtual_Appliance_10.0.0.0_IE.ova
Minimum hardware version for Resiliency Platform appliances	13 (ESXi 6.5 and above)
Resiliency Platform Data Mover (Up to hardware version 12)	Veritas_DataMover_VMware_Virtual_Appliance_10.0.0.0_IE.ova
Resiliency Platform Data Mover (From hardware version 14)	Veritas_DataMover_VMware_Virtual_Appliance_67_support_10.0.0.0_IE.ova

To download the Resiliency Platform virtual appliances:

- 1 Log in to MyVeritas portal:
<https://my.veritas.com>
- 2 Select **Licensing** tab, select the account and the entitlement ID that you want to use for downloading the Resiliency Platform virtual appliance.
- 3 In the list of products, click **Download** button next to Resiliency Platform.
- 4 Select the files that you want to download.

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

go.veritas.com/try-VRP

[Downloading the virtual appliances for Google Cloud Platform](#)

Downloading the virtual appliances for Google Cloud Platform

Below are the files which can be downloaded to be deployed on source and target data centers. The files are available in two formats: in the form of Open Virtualization Archive (OVA) files, or in the form of zip files. These .zip files contain the Virtual Hard disk (VHD) image file using which you can deploy the virtual appliances.

Table 2-2 Recovery of VMware virtual machines to Google Cloud Platform

Data Center	Component	Filenames
Source	IMS	Veritas_Resiliency_Platform_IMS_VMware_Virtual_Appliance_10.0.0.0_IE.ova
	Resiliency Platform Data Mover	Veritas_DataMover_VMware_Virtual_Appliance_10.0.0.0_IE.ova
Target	Resiliency Manager	Veritas_Resiliency_Platform_RM_VMware_Virtual_Appliance_10.0.0.0_IE.ova
	IMS	Veritas_Resiliency_Platform_IMS_VMware_Virtual_Appliance_10.0.0.0_IE.ova
	Resiliency Platform Data Mover	Veritas_DataMover_VMware_Virtual_Appliance_10.0.0.0_IE.ova

Table 2-3 Recovery of Hyper-V virtual machines to Google Cloud Platform

	Component	Filenames
Source	IMS	Veritas_Resiliency_Platform_IMS_Hyper-V_Virtual_Appliance_10.0.0.0_IE.zip
	Resiliency Platform Data Mover	Veritas_DataMover_Hyper-V_Virtual_Appliance_10.0.0.0_IE.zip

Table 2-3 Recovery of Hyper-V virtual machines to Google Cloud Platform
(continued)

	Component	Filenames
Target	Resiliency Manager	Veritas_Resiliency_Platform_RM_Hyper-V_Virtual_Appliance_10.0.0.0_IE.ova
	IMS	Veritas_Resiliency_Platform_IMS_Hyper-V_Virtual_Appliance_10.0.0.0_IE.ova
	Resiliency Platform Data Mover	Veritas_DataMover_Hyper-V_Virtual_Appliance_10.0.0.0_IE.ova

About deploying the Resiliency Platform virtual appliances

Veritas Resiliency Platform is deployed as a virtual appliance. A virtual appliance is a virtual machine image consisting of a pre-configured operating system environment with a software application installed on it. This virtual machine image can be deployed on a hypervisor. Once the Resiliency Platform virtual appliance gets deployed, you are required to configure the Resiliency Platform component through the product bootstrap.

Note: There is no sequence required for deploying and configuring the Resiliency Platform components. You can deploy and configure the components in any sequence on source as well as target data centers.

Following is the list of considerations for deploying the virtual appliances:

- For recovery to premises data center, you typically deploy and configure at least one Resiliency Manager and one Infrastructure Management Server (IMS) in the production data center and at least one Resiliency Manager and one Infrastructure Management Server (IMS) in the recovery data center.
- In case you plan to use Resiliency Platform Data Mover for recovery of your assets to premises data center, you need to deploy at least one Replication Gateway in the production data center and one Replication Gateway in the recovery data center.
- For recovery to cloud data center, you typically deploy and configure at least one Infrastructure Management Server (IMS) and one Replication Gateway in the production data center and one Resiliency Manager, one IMS, and one Replication Gateway in the recovery data center.

- The Replication Gateway on the production data center must have access to the ESX servers for the production virtual machines to be replicated. The Replication Gateway on the recovery data center must have access to the storage/ or compute or disk services of the target platform.

Resiliency Manager and IMS virtual appliance are shipped with single disk; similar to the Replication Gateway virtual appliance. While deploying the virtual appliances, it is prompted to attach a new empty disk of the required size. The RM disk size should be minimum 100 GB and IMS disk size should 40 GB.

While upgrading the virtual appliances, it will be prompted to attach the existing data disk of the previous version virtual appliances.

In the cloud environment (AWS, Azure and Google Cloud Platform), Marketplace offerings are available for upgrading the Veritas Resiliency Platform virtual appliances.

Refer to the topic **Deployment workflows** See [“Deployment workflows”](#) on page 370.

Based on the virtualization technology in your environment, choose any one of the following methods to deploy the virtual appliances in the on-premises data center:

Table 2-4 Deploying components in the on-premises data center

Virtualization technology	Steps to deploy the components
Hyper-V	See “Deploying the virtual appliance through Hyper-V Manager” on page 420.
VMware	See “Deploying the virtual appliance through VMware vSphere Client” on page 419.

Based on your cloud data center, choose any one of the following methods to deploy the virtual appliances in the cloud data center:

Table 2-5 Deploying components in the cloud data center

Cloud data center	Steps to deploy the components
Google Cloud Platform	See “Deploying the virtual appliances in Google Cloud Platform (GCP) through GCP Marketplace” on page 421. See “Deploying the virtual appliances in Google Cloud Platform using OVA files” on page 429.

Table 2-5 Deploying components in the cloud data center (*continued*)

Cloud data center	Steps to deploy the components
AWS	See “Deploying the virtual appliances in AWS through AWS Marketplace” on page 374. See “Deploying the virtual appliances in AWS using OVA files” on page 389.
vCloud Director	See “Deploying the virtual appliances in vCloud” on page 413.
Azure	See “Deploying the virtual appliances in Azure through Azure Marketplace” on page 405. See “Deploying the virtual appliances in Azure using PowerShell script” on page 396.
Azure Stack	See “Deploying the virtual appliances in Azure Stack using PowerShell script” on page 401. See “Deploy virtual appliances in Azure Stack using Azure Stack Marketplace” on page 412.
Orange Recovery Engine	See “Deploying the virtual appliances in Orange Recovery Engine” on page 415.

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

Deployment workflows

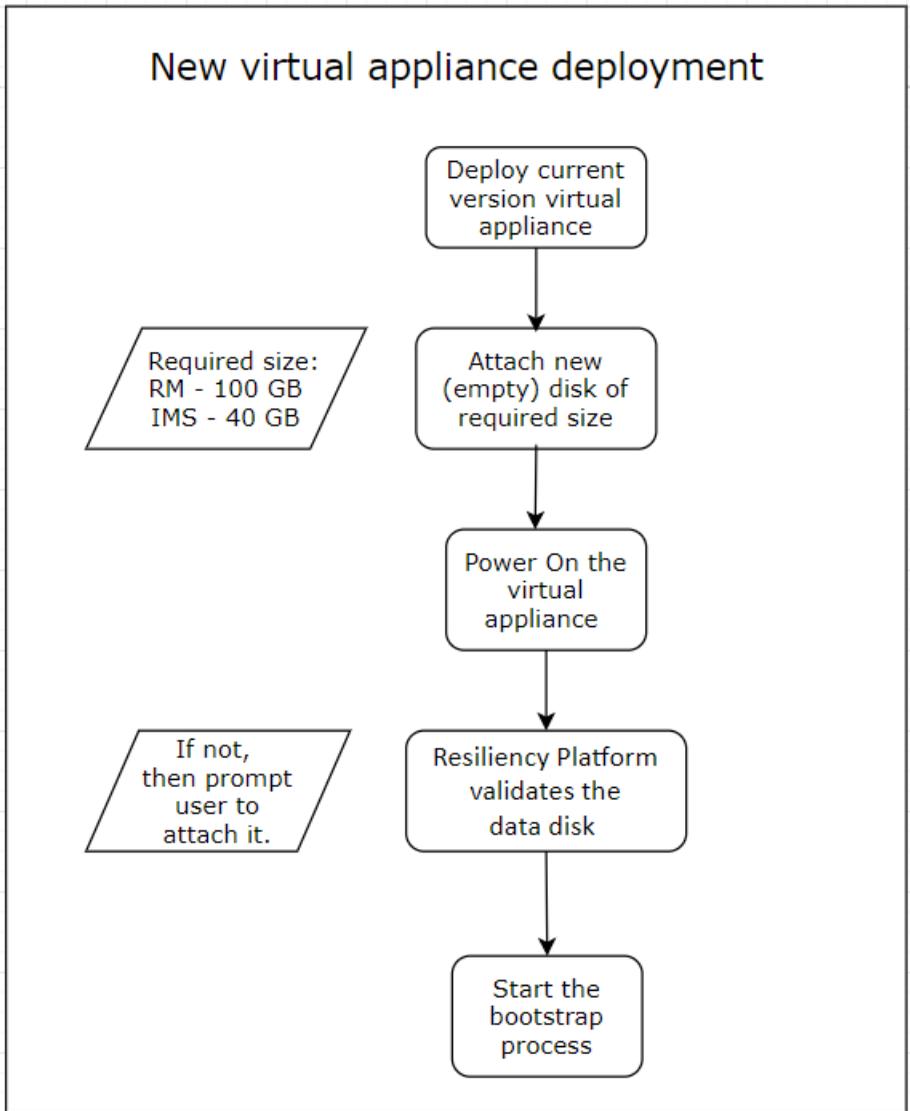
The new approach for upgrading the Resiliency Platform (Resiliency Manager and IMS) involves saving configuration to the data disk of the previous version virtual appliance and then attaching the data disk to a new, freshly deployed virtual appliance.

The Resiliency Platform virtual appliances would be created along with the OS disk. It is required to attach an existing data disk while upgrading the Resiliency Platform virtual appliances. While deploying the new virtual appliances, attach a new empty disk.

This approach is similar to that of deploying the Replication Gateway virtual appliance. It is now applicable to Resiliency Manager and IMS too. Once the disk is attached, the workflow for both the Resiliency Platform upgrade and bootstrap process is same.

Below is the deployment workflow introduced for Resiliency Manager and IMS virtual appliance:

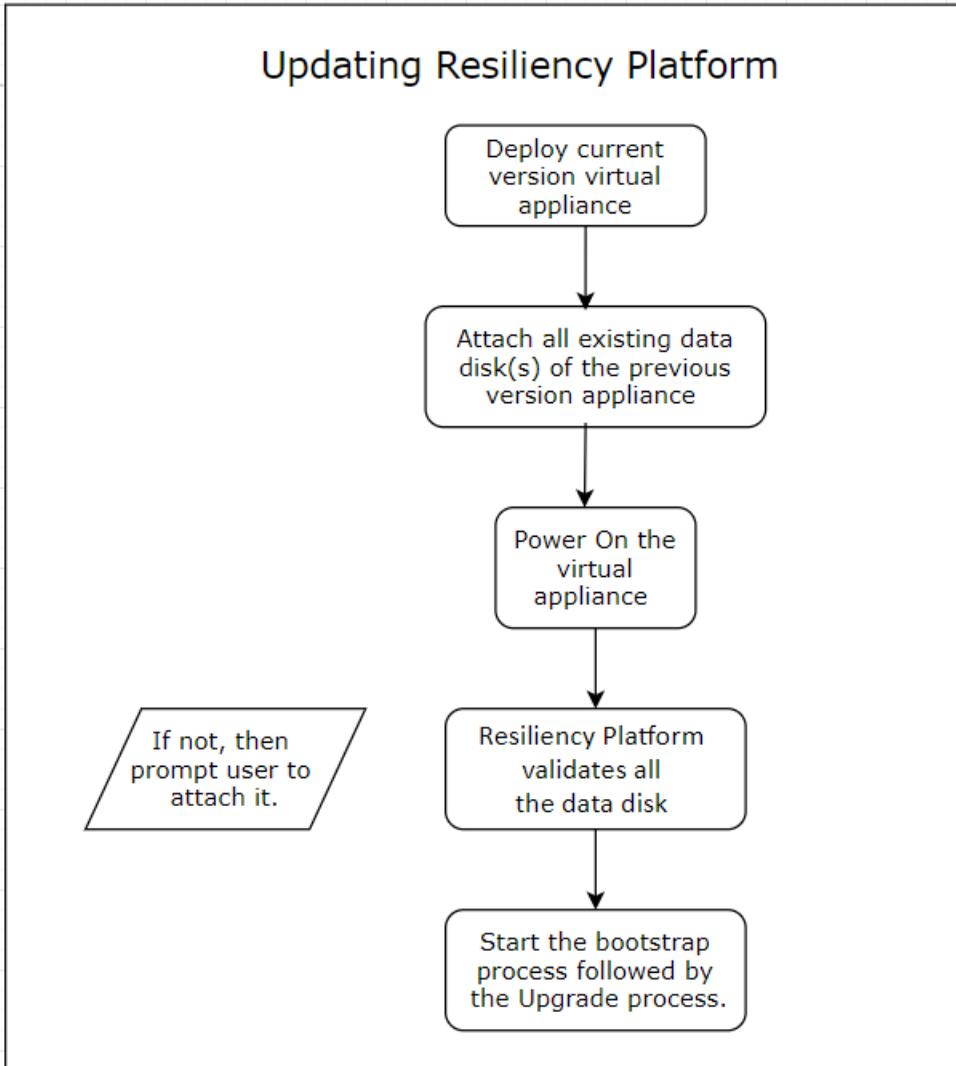
Figure 2-2 Deploying the new Resiliency Platform virtual appliances



Steps for deploying the new Resiliency Platform virtual appliances

- 1** Download the virtual appliances and deploy the current version virtual appliances.
- 2** After successfully deploying the appliances, attach an empty disk for the appliances. The required disk size for RM is minimum 100 GB and IMS is 40 GB.
- 3** Power on the virtual appliances.
- 4** Resiliency Platform validates the attached disk. Even after power on, the disk can be attached. You will be prompted to attach the disk if not attached.
- 5** Start the bootstrap process to further continue with the configuration.

Figure 2-3 Upgrading (update) the Resiliency Platform virtual appliances



Steps for upgrading Resiliency Platform virtual appliances

- 1 Download the virtual appliances and deploy the current version virtual appliances.
- 2 After successfully deploying the appliances, attach all the existing data disks of the previous version appliances.

- 3 Power on the virtual appliances.
- 4 Resiliency Platform validates all the attached data disk. Even after power on, the disk can be attached. You will be prompted to attach the data disk if not attached.
- 5 Start the bootstrap process followed by the upgrade process to further continue with the configuration.

Note: Ignore any disk or filesystem related messages given by the operating system while executing the Resiliency Platform virtual appliance upgrade or deployment bootstrap process.

Deploying the virtual appliances in AWS through AWS Marketplace

Veritas Resiliency Platform enables you to deploy the virtual appliances in AWS through AWS Marketplace using CloudFormation templates. There are six offerings available for deploying the virtual appliances using CloudFormation templates:

- **Veritas Resiliency Platform Express Install:** Installs Resiliency Manager, IMS, and Replication Gateway appliances in AWS. This template also provides options to install Veritas Data Gateway appliance in AWS.
- **Veritas Resiliency Platform Gateway Install:** Installs an additional Replication Gateway appliance in AWS.
- **Veritas Resiliency Platform Resiliency Manager Install:** Installs an additional Resiliency Manager appliance in AWS.
- **Veritas Resiliency Platform Infrastructure Management Install :** Installs the Infrastructure Management server in AWS.
- **Veritas Resiliency Platform Resiliency Manager Upgrade :** Upgrades the Resiliency Manager appliance in AWS.
- **Veritas Resiliency Platform Infrastructure Management Upgrade :** Upgrades the Infrastructure Management server in AWS.

To deploy the virtual appliances in AWS using CloudFormation templates

- 1 Prerequisites
Create Amazon Virtual Private Cloud (VPC):
- 2 Go to the **AWS Marketplace** and locate **Veritas Resiliency Platform** product.
- 3 Select the fulfillment option from the options available: **Express Install**, **Gateway Install**, **Resiliency Manager**, or **Infrastructure Management Server**.

- 4 AWS Marketplace lets you launch the selected option through CloudFormation Templates interface. You are redirected to the **Create Stack** page of AWS CloudFormation Template. The template URL is pre-populated for you. Click **Next**.
- 5 On the next page, provide the values for the input fields:
See [“Providing inputs for Resiliency Platform CloudFormation Templates”](#) on page 381.
- 6 Click **Next** and review the additional options for your stack. Select the check box in the **Capabilities** section.
- 7 Click **Next** and review the summary displayed on the **Review** page.
- 8 Click **Create** to create the instances. This step creates the EC2 instances, staging volume for Replication Gateway appliance, and required Security Groups. This step also completes the bootstrap for all the appliances.
- 9 For multiple Network Interface card (NIC) instances, every NIC created is associated with a different security group.

- 10 For security reasons, the CloudFormation template disables the SSH communication to the instances. Once the stack deployment completes, you need to enable the SSH communication to the instances for setting the admin password. This password is required for logging in to the Resiliency Manager console. A password is also required for adding the IMS, Replication Gateway appliances to Resiliency Manager.

Modify the security groups of each instance to include the inbound SSH port (TCP port 22) for the required source IP or security group.

Note: The instances go through some automatic configuration steps immediately after deployment. If you connect to the instance via SSH and if the configuration is still in-progress, wait until all steps are completed and reconnect to the instance.

- 11 Log in to the Resiliency Manager console and setup the initial infrastructure through Getting Started wizard.

Note: For **Express install**, if creation of stack fails and rollback is enabled on failure, AWS takes care of all the clean-up required in this situation. However, to enable this clean-up, you need to manually delete the Elastic Network Interface (ENI) associated with the security group that is attached with **ZipFileUploaderLambda** function.

Note: Instance Metadata Service v2 (IMDSv2) is introduced by AWS as security enhancement over IMDSv1. While using AWS CloudFormation template, there is no way to disable the IMDSv1 option. So this limits the AWS Marketplace deployment to have IMDSv2 enabled. It can be changed later using AWS CLI.

See [“Recovering VMware virtual machines to AWS”](#) on page 671.

See [“Recovering Hyper-V virtual machines to AWS”](#) on page 675.

Prerequisites for deploying the virtual appliances in AWS

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

- Follow the documentation of AWS to create the required security groups. make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open for communication.

If you deploy Resiliency Platform components through AWS Marketplace, the required security groups are automatically created.

- Create a role named *vmimport* and grant the permissions required to import an image to the role. Follow the documentation of AWS to know about the permissions required to import an image.

The *vmimport* role should have the following KMS service permissions (along with the permissions mentioned in the AWS documentation):

- "kms:ReEncrypt"
 - "kms:GenerateDataKeyWithoutPlaintext"
 - "kms:DescribeKey"
 - "kms:CreateGrant"
 - "kms:Decrypt"
- Create Individual roles for Resiliency Manager, IMS, and Replication Gateway with certain permissions. These roles are used for authenticating the operations performed by the Resiliency Platform components in AWS.

See [“Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway”](#) on page 389.

If you deploy Resiliency Platform through AWS Marketplace, then the required role are automatically created through AWS CloudFormation template that the marketplace deployment uses.

- Ensure that there is direct communication between the premise network and the AWS network. It is recommended to use VPN for AWS environment.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. You may choose to restrict incoming internet access on these virtual appliances.
- Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.
- For significant cost benefits, it is recommended to buy Amazon EC2 reserved instances for the Resiliency Platform virtual appliances as the virtual appliances will be running continuously. While buying the reserved instances, it is also recommended to select the availability zones where the virtual appliances are to be deployed; this will ensure reserved capacity. It is important to choose the reserved instance types that matches the virtual appliances' instance types.

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

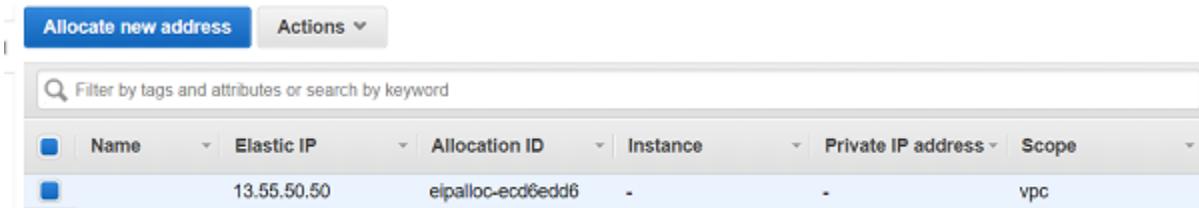
Configuring Amazon VPC for deployment using CloudFormation Templates

You need to configure Amazon Virtual Private Cloud (VPC) for deploying the Veritas Resiliency Platform components in AWS using CloudFormation Templates (CFT). For information on deploying the Veritas Resiliency Platform components in AWS using CFT:

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

To configure Amazon VPC

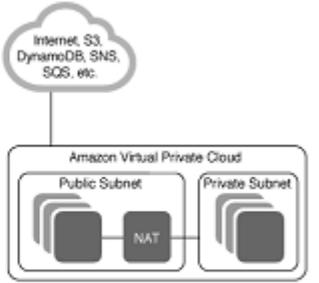
- 1 Access the VPC service in the AWS region where you want to deploy Resiliency Platform components, and go to the **Elastic IPs** tab.
- 2 Click on **Allocate new address** to allocate an Elastic IP. This Elastic IP will be required for the NAT gateway while creating the VPC.



- 3 In the **VPC Dashboard** tab, click on **Start VPC Wizard**.

- 4 In the **Select a VPC Configuration** step, select the **VPC with Public and Private Subnets** option.

Step 1: Select a VPC Configuration

<p>VPC with a Single Public Subnet</p>	<p>In addition to containing a public subnet, this configuration adds a private subnet whose instances are not addressable from the Internet. Instances in the private subnet can establish outbound connections to the Internet via the public subnet using Network Address Translation (NAT).</p> <p>Creates:</p> <p>A /16 network with two /24 subnets. Public subnet instances use Elastic IPs to access the Internet. Private subnet instances access the Internet via Network Address Translation (NAT). (Hourly charges for NAT devices apply.)</p> <p style="text-align: right;">Select</p>	
<p>VPC with Public and Private Subnets</p>		
<p>VPC with Public and Private Subnets and Hardware VPN Access</p>		
<p>VPC with a Private Subnet Only and Hardware VPN Access</p>		

- 5 In the **VPC with Public and Private Subnets** page, do the following:
 - Choose an appropriate CIDR block for the VPC and give it the desired name. Repeat the same steps for the public and private subnets.
 - For the NAT Gateway, select the Elastic IP that was allocated in step 2.
 - Ensure that the **Enable DNS hostnames** option is set to **Yes**.

Step 2: VPC with Public and Private Subnets

IPv4 CIDR block:* (65531 IP addresses available)

IPv6 CIDR block: No IPv6 CIDR Block
 Amazon provided IPv6 CIDR block

VPC name:

Public subnet's IPv4 CIDR:* (251 IP addresses available)

Availability Zone:*

Public subnet name:

Private subnet's IPv4 CIDR:* (251 IP addresses available)

Availability Zone:*

Private subnet name:

You can add more subnets after AWS creates the VPC.

Specify the details of your NAT gateway ([NAT gateway rates apply](#)).

Elastic IP Allocation ID:*

Service endpoints

Enable DNS hostnames:* Yes No

Hardware tenancy:*

6 Finish the wizard to create the VPC.

Providing inputs for Resiliency Platform CloudFormation Templates

You need to provide inputs for creating instances using CloudFormation templates (CFT). Some of the fields get auto populated with the default value, you can change the values if required. For rest of the parameters, you need to enter a valid value.

Table 2-6 EC2 Instance Configuration

Field	Description
Resiliency Manager Instance Name	Enter the name of the Resiliency Manager instance.
IMS Instance Name	Enter the name of the IMS instance.
Replication Gateway Instance Name	Enter the name of the Replication Gateway instance.
EC2 Instance Type	This field is auto-populated with the best fit instance type. You can change the default value, if required.
Key Pair for SSH access	Select the appropriate SSH key pair. While selecting the SSH key pair, ensure that you already have access to the pair. This is critical because SSH key pair is the only way to log in to the system and you cannot change the SSH key pair after the EC2 instance gets created.
Create 'ImportSnapshotRole' AWS IAM Role for VRP?	Resiliency Platform requires ImportSnapshotRole for recovery of assets to AWS. Select Yes to create the role. Make sure that you have iam: CreateRole permission. Select No if you already have ImportSnapshotRole role.
Staging Disk size for the Replication Gateway	Enter the size of the staging disk for Replication Gateway in the unit of GB. The minimum supported disk size of 50GB lets you protect up to 8 virtual machines and to protect more virtual machines an additional staging storage of 6 GB per virtual machine is required.

Table 2-6 EC2 Instance Configuration (*continued*)

Field	Description
Data Volume Type	Select the appropriate Volume Type for the data volume. This field has default value selected as 'gp2' volume type. You can change and select another as per your requirement.
Resiliency Manager Data volume IOPS (Required only when selected Volume Type is IO1 from IO1 and IO2)	This value will be used when selected Volume Type is IO1 only .IOPS value must be between 100 to 64000.
Infrastructure Management Server Data volume IOPS (Required only when selected Volume Type is IO1 from IO1 and IO2)	This value will be used when selected Volume Type is IO1 only. IOPS value must be between 100 to 64000.
Replication Gateway Staging Disk IOPS (Required only when selected Volume Type is IO1 from IO1 and IO2)	<p>This value will be used when selected DataVolumeType is IO1. The maximum ratio of provisioned IOPS to requested volume size (in GiB) is 50:1. For example, a 100 GiB volume can be provisioned with up to 5,000 IOPS. If appropriate IOPS value is not selected , the deployment fails.</p> <p>This value will be used when selected DataVolumeType is IO2. The maximum ratio of provisioned IOPS to requested volume size (in GiB) is 500:1. For example, a 100 GiB volume can be provisioned with up to 50,000 IOPS. If appropriate IOPS value is not selected , the deployment fails.</p>
KMS Key ARN for Staging Disk encryption	If you want to encrypt staging disks attached to the gateways using the KMS key, then select ARN of the respective KMS from the KMS Key drop-down. Else keep the default value as 'Not Encrypted'.

Table 2-7 Resiliency Manager Network Configuration

Field	Description
Network Interface to be used for communication with other Resiliency Managers	<p>Select the network interface that can be used to communicate with other Resiliency Managers in the resiliency domain.</p> <p>NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.</p>
Network Interface to be used for communication with Infrastructure Management Servers	<p>Select the network interface that can be used to communicate with Infrastructure Management Servers.</p> <p>NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.</p>
Network Interfaces to be used for accessing the User Interface	<p>Select the network interfaces that can be used to access the Resiliency Manager web user interface.</p> <p>NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.</p>
Network Interface to be used as the default gateway	Select the network interface that can be used by default for outgoing communication.
Resiliency Manager eth0 Subnet	Select the subnet which can be connected to the eth0 network interface.
Resiliency Manager eth1 Subnet	<p>Select the subnet which can be connected to the eth1 network interface.</p> <p>This input is only relevant if you have selected eth1 for one of the communication roles. You should provide a valid subnet always.</p>
Is the eth0 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth0 and if you wish to use NAT with eth0, select Yes.
Resiliency Manager eth0 NAT Hostname (Optional)	Provide the NAT hostname for eth0 if applicable

Table 2-7 Resiliency Manager Network Configuration (*continued*)

Field	Description
Resiliency Manager eth0 NAT IP (Optional)	Provide the NAT IP for eth0 if applicable
Is the eth1 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth1 and if you wish to use NAT with eth1, select Yes.
Resiliency Manager eth1 NAT Hostname (Optional)	Provide the NAT hostname for eth1 if applicable
Resiliency Manager eth1 NAT IP (Optional)	Provide the NAT IP for eth1 if applicable

Table 2-8 Infrastructure Management Server Network Configuration

Field	Description
Network Interface to be used for communication with Resiliency Managers	Select the network interface that can be used to communicate with Resiliency Managers. NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.
Network Interface to be used for communication with Replication Gateways	Select the network interface that can be used to communicate with Replication Gateways. NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.
Network Interface to be used as the default gateway	Select the network interface that can be used by default for outgoing communication.
IMS eth0 Subnet	Select the subnet which can be connected to the eth0 network interface.
IMS eth1 Subnet	Select the subnet which can be connected to the eth1 network interface. This input is only relevant if you have selected eth1 for one of the communication roles. You should provide a valid subnet always.

Table 2-8 Infrastructure Management Server Network Configuration
(continued)

Field	Description
Is the eth0 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth0 and if you wish to use NAT with eth0, select Yes.
IMS eth0 NAT Hostname (Optional)	Provide the NAT hostname for eth0 if applicable
IMS eth0 NAT IP (Optional)	Provide the NAT IP for eth0 if applicable
Is the eth1 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth1 and if you wish to use NAT with eth1, select Yes.
IMS eth1 NAT Hostname (Optional)	Provide the NAT hostname for eth1 if applicable
IMS eth1 NAT IP (Optional)	Provide the NAT IP for eth1 if applicable

Table 2-9 Replication Gateway Network Configuration

Field	Description
Network Interface to be used for communication with the Infrastructure Management Server	Select the network interface that can be used to communicate with the Infrastructure Management Server. NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.
Network Interface to be used for communication with peer gateways	Select the network interface that can be used to communicate with peer gateways. NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.

Table 2-9 Replication Gateway Network Configuration (*continued*)

Field	Description
Network Interface to be used for communication with workload Virtual Machines	<p>Select the network interfaces that can be used for communication with the workload Virtual Machines.</p> <p>NAT settings are applicable for this role. If you want to use this network interface in a NAT environment, provide NAT details in the form.</p>
Network Interface to be used as the default gateway	Select the network interface that can be used by default for outgoing communication.
Replication Gateway eth0 Subnet	Select the subnet which can be connected to the eth0 network interface.
Replication Gateway eth1 Subnet	<p>Select the subnet which can be connected to the eth1 network interface.</p> <p>This input is only relevant if you have selected eth1 for one of the communication roles. You should provide a valid subnet always.</p>
Replication Gateway eth2 Subnet	<p>Select the subnet which can be connected to the eth2 network interface.</p> <p>This input is only relevant if you have selected eth2 for one of the communication roles. You should provide a valid subnet always.</p>
Is the eth0 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth0 and if you wish to use NAT with eth0, select Yes.
Replication Gateway eth0 NAT Hostname (Optional)	Provide the NAT hostname for eth0 if applicable
Replication Gateway eth0 NAT IP (Optional)	Provide the NAT IP for eth0 if applicable
Is the eth1 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth1 and if you wish to use NAT with eth1, select Yes.
Replication Gateway eth1 NAT Hostname (Optional)	Provide the NAT hostname for eth1 if applicable

Table 2-9 Replication Gateway Network Configuration (*continued*)

Field	Description
Replication Gateway eth1 NAT IP (Optional)	Provide the NAT hostname for eth1 if applicable
Is the eth2 Network Interface behind NAT?	If NAT settings are applicable for the communications you have selected for eth2 and if you wish to use NAT with eth2, select Yes.
Replication Gateway eth2 NAT Hostname (Optional)	Provide the NAT hostname for eth2 if applicable
Replication Gateway eth2 NAT IP (Optional)	Provide the NAT IP for eth2 if applicable

Table 2-10 Common Network Configuration

Field	Description
VPC ID	Ensure the following while selecting the VPC: <ul style="list-style-type: none"> ■ Outgoing internet access is enabled from at least one private subnet from the VPC. ■ VPC has a VPN configured with the network at the on- premises data center.

Table 2-11 Resiliency Platform Installation Parameters

Field	Description
NTP Server	FQDN or IP address of the NTP server to be used for the instances. In case of multiple values, enter the space-separated values.
TimeZone	Select the timezone for the instances.

Table 2-12 Data Gateway Deployment Information

Field	Description
Data Gateway deployment bucket name	Make sure that the deployment bucket pre-exists and is created in the local region where the CFT is being deployed.
Deploy Veritas Data Gateway?	Optional parameter You need to deploy Data Gateway only if you want to use Object Storage for replication.

Table 2-12 Data Gateway Deployment Information (*continued*)

Field	Description
Does the Data Gateway bucket already exist?	Select Yes if the Data Gateway already exists. If you want to re-use an existing Data Gateway Bucket, that bucket should also be local to the region in which the CFT is being deployed.
SNS Topic Protocol	Select a protocol from the list.
SNS Topic Endpoint	Provide an endpoint that is appropriate for the selected SNS topic protocol.

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Uninstalling Resiliency Platform components when deployed through AWS Marketplace

When you deploy Resiliency Platform components through AWS Marketplace, a stack gets created in AWS environment for each offering that you use. For example, if you use Express install, one stack gets created and all the Resiliency Platform components are deployed in that stack. If you use Resiliency Manager offering to install additional Resiliency Manager, another stack gets created. If you use Replication Gateway offering to install additional Replication Gateway, one more stack gets created.

If at a later point of time, you want to uninstall all the Resiliency Platform components that were deployed through AWS Marketplace, you need to delete all these stacks individually. Once you delete the stacks, all the resources that were created during deployment automatically get deleted.

To delete a stack in AWS

- 1 From the list of stacks in the AWS CloudFormation console, select the stack that you want to delete.
- 2 Click **Actions > Delete stack**.
- 3 When prompted, confirm that you want to delete the stack.

If you want to delete a stack created for **Express install** within 6 hours of creating the stack, you need to manually delete the Elastic Network Interface (ENI) associated with the security group that is attached with **ZipFileUploaderLambda** function.

Rest of the stack gets deleted automatically when you perform **Delete stack** operation.

Note: The NICs attached to Resiliency Platform virtual appliances that are deployed through the Marketplace will not be deleted automatically when the instances are deleted. They would need to be cleaned up manually after deletion.

Deploying the virtual appliances in AWS using OVA files

To know about virtual appliance deployment in Veritas Resiliency Platform:

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

Table 2-13 Overview of deployment process in AWS

Step	Action	Description
1	Ensure that the prerequisites for deploying virtual appliances in AWS are met.	See “Prerequisites for deploying the virtual appliances in AWS” on page 376.
2	Upload the OVA files to Amazon S3	See “Uploading the OVA file using web-based method” on page 393. See “Uploading the OVA file using command-line method” on page 394.
3	Create AMI using EC2	See “Creating Amazon Machine Image” on page 394.
4	Launch the instances of virtual appliances to deploy Resiliency Manager, Infrastructure Manager (IMS), and Replication Gateway	See “Launching the instances of virtual appliances” on page 395.

Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway

Following are the permissions required for the roles that you need to create for Resiliency Manager, IMS, and Replication Gateway for recovery to AWS data center.

Table 2-14 Permissions required for role for Resiliency Manager

Service name	Permission
ec2	ec2:DescribeVpcs ec2:DescribeAvailabilityZones
S3	s3:ListBucket

Table 2-15 Permissions required for role for IMS

Service name	Permission
ec2	

Table 2-15 Permissions required for role for IMS (*continued*)

Service name	Permission
	ec2:RunInstances
	ec2:CreateTags
	ec2:StartInstances
	ec2:DescribeInstanceStatus
	ec2:StopInstances
	ec2:TerminateInstances
	ec2:RegisterImage
	ec2:DescribeImageAttribute
	ec2:DescribeInstanceAttribute
	ec2:DeregisterImage
	ec2>DeleteVolume
	ec2:CreateVolume
	ec2:AttachVolume
	ec2:DetachVolume
	ec2:CreateSnapshot
	ec2>DeleteSnapshot
	ec2:ImportSnapshot
	ec2:DescribeImportSnapshotTasks
	ec2:CreateImage
	ec2:CreateNetworkInterface
	ec2>DeleteNetworkInterface
	ec2:AttachNetworkInterface
	ec2:ModifyNetworkInterfaceAttribute
	ec2:DescribeVpcs
	ec2:DescribeSubnets
	ec2:DescribeImages
	ec2:DescribeSecurityGroups
	ec2:DescribeNetworkInterfaces
	ec2:DescribeVolumes
	ec2:DescribeAvailabilityZones

Table 2-15 Permissions required for role for IMS (*continued*)

Service name	Permission
	ec2:DescribeSnapshots ec2:DescribeInstances ec2:GetEbsDefaultKmsKeyId ec2:GetEbsEncryptionByDefault ec2:DescribeInstanceTypes
execute-api	execute-api:Invoke
S3	s3:ListBucket s3:PutObject s3>DeleteObject s3:GetObject s3:GetBucketLocation
KMS	kms:ListAliases kms:ListKeys kms:DescribeKey

Table 2-16 Permissions required for role for Replication Gateway

Service name	Permission
execute-api	execute-api:Invoke
S3	s3:PutObject s3:GetObject s3:ListBucket

Uploading the OVA file using web-based method

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

To upload the OVA file using web-based method

- 1 Log in to the AWS console and go to **Services**.
- 2 Go to **S3**, and click **Create a bucket**.
- 3 Enter a name for the bucket and select the appropriate region. Click **Create**.

- 4 Once the bucket gets created, open the bucket and click **Upload**. Click **Add files** and then select the OVA file from your local disk.
- 5 Click **Start Upload**.

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Uploading the OVA file using command-line method

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

To upload the OVA file using command-line method

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

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

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

If you do not use the region option of the command, the bucket is created in the region specified in your configuration file.

- 3 Upload the OVA file by running the following command:

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

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

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Creating Amazon Machine Image

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

To create Amazon Machine Image

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

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

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

- 3 Run the following command to create an AMI:

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

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

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

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

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

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Launching the instances of virtual appliances

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

Instance MetaData Service v2 (IMDSv2) is introduced by AWS as security enhancement over IMDSv1. Instance MetaData Service v2 (IMDSv2) is introduced by AWS as security enhancement over IMDSv1. From version 4.0 of Resiliency Platform, while configuring resiliency group for disaster recovery, the wizard has an option to specify metadata access using **Enforce IMDSv2** option per virtual machine. If this option is true, the virtual machine when migrated to AWS, should use only IMDSv2 mechanism.

To launch the instances of virtual appliances

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

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

In the **Select an existing key pair or create a new key pair** wizard, you can choose an existing key pair, or create a new one. If you create a new key pair, ensure to click the **Download key pair** button and download. You will need the private key from this key pair to login as admin user for completing the bootstrap.

See [“Deploying the virtual appliances in AWS through AWS Marketplace”](#) on page 374.

Deploying the virtual appliances in Azure using PowerShell script

To know about virtual appliance deployment in Veritas Resiliency Platform:

You can deploy the Resiliency Platform virtual appliances using the following two files provided by Veritas along with the product files:

- A PowerShell script that handles the entire deployment of virtual appliances in Azure environment.
- A text file that has all the parameters required for deployment in Azure environment. You need to update the values assigned to the parameters in the text file. These values are used by the PowerShell script while deploying the virtual appliances.

To deploy the virtual appliances in Azure

- 1** Ensure that the prerequisites for deploying virtual appliances in Azure are met. See [“Prerequisites for deploying the virtual appliances in Azure and Azure Stack”](#) on page 403.
- 2** Log in to the Azure portal and create a general purpose storage account (and not account type as blob storage). Create a container with private access type under this storage account. Follow the documentation of Azure to create the storage account and container.
- 3** Create a static network interface for the virtual appliance. A static network interface in Azure ensures that the IP of the appliance does not change after reboot.
- 4** Download or copy the Azure deployment files on your local Windows system.

- 5 Update the values of all the parameters in the text file `VRPVSADeployInputs.txt` according to your environment.

Virtual Machine Appliance	Compatible Sizes	Recommended Sizes
Resiliency Manager	Standard_DS5_v2	Standard_D8s_v3
	Standard_DS13_v2	Standard_DS5_v2
	Standard_E8s_v3	
	Standard_D8s_v3	
	Standard_E8_v3	
	Standard_F16s_v2	
	Standard_F16s	
	Standard_A8m_v2	
Infrastructure Management Server	Standard_D13	
	Standard_F8s	Standard_F8s
	Standard_F8s_v2	Standard_F8s_v2
	Standard_D8s_v3	Standard_A8_v2
	Standard_DS4_v2	
	Standard_D4_v2	
	Standard_DS4	
	Standard_D4	
Replication Gateway	Standard_A8_v2	
	Standard_A4	
	Standard_F8s	Standard_F8s
	Standard_F8s_v2	Standard_F8s_v2
	Standard_D8s_v3	Standard_A8_v2
	Standard_DS4_v2	
	Standard_D4_v2	
	Standard_DS4	
Standard_D4		
	Standard_A8_v2	
	Standard_A4	

It is recommended that the Replication Gateway should be of virtual machine size which supports ultra disk and premium storage disk type.

- 6 Open the PowerShell command prompt and run the following command to log in to Azure:

```
Add-AzureRmAccount
```

- 7 Enter the absolute path and filename and then press the **Enter** key to run the Powershell script.

Prerequisites for deploying the virtual appliances in Azure and Azure Stack

Following are the prerequisites for deploying the Resiliency Platform virtual appliances in Azure. These prerequisites are applicable to deploy virtual appliances in Azure Stack too:

- Follow the documentation of Azure to create the required security groups. make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open.
- Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.
- Ensure that there is direct communication between the premise network and the Azure network. It is recommended to use VPN for Azure environment.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. You may choose to restrict incoming internet access on these virtual appliances.
- Ensure that PowerShell is installed on the Windows system from which you plan to run the Azure deployment script.
- Ensure that Azure PowerShell module is installed on the system.
See [“Installing Azure PowerShell module”](#) on page 399.

See [“Deploying the virtual appliances in Azure using PowerShell script”](#) on page 396.

Installing Azure PowerShell module

Installing Azure PowerShell from the PowerShell Gallery is the preferred method of installation.

To install Azure PowerShell module

- 1 You should have PowerShellGet module installed on your system.

<https://www.microsoft.com/en-us/download/details.aspx?id=51451>

- 2 You can verify if PowerShellGet module is properly installed on your system by executing the following command:

```
Get-Module PowerShellGet -list | Select-Object Name,Version,Path
```

If PowerShellGet module is properly installed on your system, you get the output of the above command similar to the following:

```
PowerShellGet 1.0.0.1 C:\Program  
Files\WindowsPowerShell\Modules\PowerShellGet\1.0.0.1\PowerShellGet.psd1
```

- 3 Once PowerShellGet module gets installed on your system, install the **Azure Resource Manager** modules by running the following command as an administrator:

- `Install-Module -Name Az`

- 4 You can check if Azure Resource Manager is properly installed by running any one of the following commands:

- `Add-AzAccount`

- `Get-AzResourceGroup`

See “[Prerequisites for deploying the virtual appliances in Azure and Azure Stack](#)” on page 403.

See “[Deploying the virtual appliances in AWS through AWS Marketplace](#)” on page 374.

Constraints when you deploy the virtual appliances in Azure

Following are a few limitations that apply when you deploy the Resiliency Platform virtual appliances in Azure:

- To prevent any security vulnerabilities, `Waagent` service is stopped 30 minutes after the bootstrapping completes on all the Resiliency Platform appliances. If required, you can start the service for next 30 minutes by using the following Klish command:

```
azure-waagent-service start
```

- You must not use `run-command` option from Azure portal on any of the Resiliency Platform appliances. Running `run-command` on the appliances may cause some functionality related issues for the product.

- Extensions are disabled on all the Resiliency Platform appliances.

Deploying the virtual appliances in Azure Stack using PowerShell script

To know about virtual appliances in Veritas Resiliency Platform:

You can deploy Resiliency Platform virtual appliances using the following two files provided by Veritas along with the product files:

- A PowerShell script that handles the entire deployment of virtual appliances in Azure Stack environment.
- A text file that has all the parameters required for deployment in Azure Stack environment. You need to update the values assigned to the parameters in the text file. These values are used by the PowerShell script while deploying the virtual appliances

To deploy the virtual appliances in Azure

- 1** Ensure that the prerequisites for deploying virtual appliances in Azure Stack are met.

See [“Prerequisites for deploying the virtual appliances in Azure and Azure Stack”](#) on page 403.
- 2** Log in to the Azure Stack portal and create a general purpose storage account (and not account type as blob storage). Create a container with private access type under this storage account. Follow the documentation of Azure to create the storage account and container.
- 3** Create a static network interface for the virtual appliance. A static network interface in Azure ensures that the IP of the appliance does not change after reboot.
- 4** Download or copy the Azure deployment files on your local Windows system.

- 5 Update the values of all the parameters in the text file `VRPVSADeployInputs.txt` according to your environment.

Virtual Machine Appliance	Compatible Sizes	Recommended Sizes
Resiliency Manager	Standard_DS5_v2	Standard_D8s_v3
	Standard_DS13_v2	Standard_DS5_v2
	Standard_E8s_v3	
	Standard_D8s_v3	
	Standard_E8_v3	
	Standard_F16s_v2	
	Standard_F16s	
	Standard_A8m_v2	
Infrastructure Management Server	Standard_D13	
	Standard_F8s	Standard_F8s
	Standard_F8s_v2	Standard_F8s_v2
	Standard_D8s_v3	Standard_A8_v2
	Standard_DS4_v2	
	Standard_D4_v2	
	Standard_DS4	
	Standard_D4	
Replication Gateway	Standard_AB_v2	
	Standard_A4	
	Standard_F8s	Standard_F8s
	Standard_F8s_v2	Standard_F8s_v2
	Standard_D8s_v3	Standard_A8_v2
	Standard_DS4_v2	
	Standard_D4_v2	
	Standard_DS4	
Standard_D4		
Standard_A8_v2		
Standard_A4		

It is recommended that the Replication Gateway should be of virtual machine size which support premium storage type.

- 6 Open the PowerShell command prompt and run the following command to log in to Azure:

To add a new Azure environment:

```
Add-AzEnvironment -Name <env_name> -ArmEndpoint <endpoint_url>
```

To connect to Azure environment:

```
Connect-AzAccount -Environment <env_name> -SubscriptionId  
<subscription_id>
```

- 7 Enter the absolute path and filename and then press the **Enter** key to run the PowerShell script.

Prerequisites for deploying the virtual appliances in Azure and Azure Stack

Following are the prerequisites for deploying the Resiliency Platform virtual appliances in Azure. These prerequisites are applicable to deploy virtual appliances in Azure Stack too:

- Follow the documentation of Azure to create the required security groups. make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open.
- Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.
- Ensure that there is direct communication between the premise network and the Azure network. It is recommended to use VPN for Azure environment.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. You may choose to restrict incoming internet access on these virtual appliances.
- Ensure that PowerShell is installed on the Windows system from which you plan to run the Azure deployment script.
- Ensure that Azure PowerShell module is installed on the system. See [“Installing Azure PowerShell module”](#) on page 399.

See [“Deploying the virtual appliances in Azure using PowerShell script”](#) on page 396.

Installing Azure Stack PowerShell module

Installation of PowerShell module depends on Azure stack version. You can refer Microsoft documentation. The below mentioned steps are for Azure Stack version 1908:

- Install Windows PowerShell 5.1.
- Install `PowerShellGet`: You need access to the PowerShell Gallery. The gallery is the central repository for PowerShell content. The `PowerShellGet` module contains cmdlets for discovering, installing, updating, and publishing PowerShell artifacts.
 - To install a package from the Gallery either execute the `Install-Module` or `Install-Script` cmdlet, depending on the package type.
Run: `Install-Module -Name PowershellGet -RequiredVersion 2.2.1`
 - Validate the PowerShell Gallery accessibility using the below commands:
`Import-Module -Name PowerShellGet -ErrorAction Stop`
`Import-Module -Name PackageManagement -ErrorAction Stop`
`Get-PSRepository -Name "PSGallery"`
 - If the repository isn't registered, open an elevated PowerShell session and run the following command:
`Register-PSRepository -Default Set-PSRepository -Name "PSGallery" -InstallationPolicy Trusted`

Install Azure Stack PowerShell

To install the Azure Stack PowerShell, perform the below steps:

- Install the `AzureRM.BootStrapper` module. Select **Yes** when prompted to install NuGet.
`Install-Module -Name AzureRM.BootStrapper`
- Install and import the API Version Profile required by Azure Stack into the current PowerShell session. Required version will vary with respect to Azure Stack Version
`Use-AzureRmProfile -Profile 2019-03-01-hybrid -Force Install-Module -Name AzureStack -RequiredVersion 1.7.2`

Confirm the installation

You can confirm the Azure Stack PowerShell installation by executing the below commands:

```
Get-Module -Name "Azure*" -ListAvailable  
Get-Module -Name "Azs*" -ListAvailable
```

See [“Deploying the virtual appliances in Azure Stack using PowerShell script”](#) on page 401.

Constraints when you deploy the virtual appliances in Azure Stack

Following are a few limitations that apply when you deploy the Resiliency Platform virtual appliances in Azure Stack:

- To prevent any security vulnerabilities, `Waagent` service is stopped 30 minutes after the bootstrapping completes on all the Resiliency Platform appliances. If required, you can start the service for next 30 minutes by using the following Klish command:

```
azure-waagent-service start
```
- You must not use `run-command` option from Azure portal on any of the Resiliency Platform appliances. Running `run-command` on the appliances may cause some functionality related issues for the product
- Extensions are disabled on all the Resiliency Platform appliances.

See [“Deploying the virtual appliances in Azure Stack using PowerShell script”](#) on page 401.

Deploying the virtual appliances in Azure through Azure Marketplace

To know about virtual appliance deployment in Veritas Resiliency Platform:

A few constraints are applied when you deploy Resiliency Platform in Azure:

See [“Constraints when you deploy the virtual appliances in Azure”](#) on page 400.

Veritas Resiliency Platform enables you to deploy the virtual appliances in Azure through Azure Marketplace using Azure Resource Manager (ARM) templates. To deploy the appliances search for the offering **Veritas Resiliency Platform on the Azure marketplace..**

The offering provide choice to create a set of Resiliency Manager, Infrastructure Management server (IMS) and Replication Gateway or to deploy these appliances individually.

To deploy the virtual appliances in Azure through Azure Marketplace

1 Prerequisites

Ensure that the prerequisites are met:

See [“Prerequisites for deploying the virtual appliances in Azure and Azure Stack”](#) on page 403.

2 Go to the Azure Marketplace and locate **Veritas Resiliency Platform** offering.

3 Select the offering. Azure Marketplace lets you launch the selected offering.

4 You are redirected to the **Deployment** page of the offering. Click **Create** to initiate the process.

5 On the next page, provide the values for the input fields:

See [“Providing inputs for deploying virtual appliances through Azure Marketplace”](#) on page 406.

6 Click **OK** and review the summary displayed in the **Summary** section.

7 Click **OK** to view and accept the terms and conditions.

8 Click **Create** to create the instances:

- This step creates selected virtual appliance(s) and completes the bootstrap of appliances with the provided inputs. It also creates the required network security groups and network interfaces. An additional staging disk is also created during Replication Gateway deployment.

9 For security reasons, the template disables the SSH communication to the instances. Once the deployment completes, you need to enable the SSH communication to the instances. The password is required for logging in to the Resiliency Manager console. The password is also required for adding the IMS, Replication Gateway to Resiliency Manager.

Modify the security group of each instance to include the inbound SSH port (TCP port 22) for the required source IP or security group.

10 Log in to the Resiliency Manager console and setup the initial infrastructure through Getting Started wizard.

Providing inputs for deploying virtual appliances through Azure Marketplace

You need to provide inputs for creating instances using Azure Resource Manager (ARM) templates. Some of the fields get auto populated with the default value, you can change the values if required. For rest of the parameters, you need to enter a valid value.

To create the appliance images, you need to provide inputs in following three sections:

- Basics [Table 1-17](#)
- Advanced [Table 1-18](#)
- Network [Table 1-19](#)

Figure 2-4 Basic settings for Azure deployment

The screenshot displays the 'Basics' step of an Azure deployment configuration. At the top, there are four numbered tabs: 1 Basics (selected), 2 Advanced, 3 Resiliency Manager Network, and 4 Infrastructure Management Network. Below the tabs, the 'Project details' section includes a heading and a sub-heading: 'Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.' This section contains two dropdown menus: 'Subscription *' with the value 'VRP-team Subscription' and 'Resource group *' with the value '(New) VRP'. A 'Create new' link is visible below the resource group dropdown. The 'Instance details' section contains a 'Region *' dropdown menu with the value 'East US'. The 'Resiliency Platform Bootstrap Inputs' section contains a 'Select deployment type' dropdown menu with the value 'All Appliances (fresh install)'. A grey information banner at the bottom of the form contains an information icon and the text: 'It will deploy a fresh Resiliency Manager, an Infrastructure Management Server (IMS) and a Replication Gateway'. At the very bottom, there are two buttons: '< Previous' and 'Next'.

Table 2-17 Basic settings for Azure deployment

Input field	Description
Select Deployment Type	Select the appropriate deployment type from the given options. By default, all the appliances are selected which will deploy a Resiliency Manager, an IMS, a Replication Gateway, Resiliency Manager upgrade and IMS upgrade appliance. You can select other option as per the requirement.
Password for admin user	Set the password for admin user. The admin user and password is later used for configuring the appliances.
Confirm password	Provide same password for confirmation.
NTP Servers	Resiliency Platform requires NTP servers for synchronizing time across all the appliances. Specify FQDN or IP addresses of one or more (space separated) NTP servers. It is recommended to use 3 or more (an odd number) NTP servers.
Timezone	Timezone to be set for all the Resiliency Platform appliances.
Subscription	Select the subscription of Azure account, to be used for deploying the virtual appliances.
Resource group	Specify the name of an existing resource group or a new resource group to be created.
Location	Select the location where you want to create the appliances.

Figure 2-5 Advanced settings for Azure deployment

The screenshot shows the 'Advanced' settings page in the Azure portal. At the top, there are navigation tabs: 'Basics' (selected), 'Advanced', 'Resiliency Manager Network', 'Infrastructure Management Network', 'Replication Gateway Network Settings', and 'Review + create'. The 'Advanced' tab is active, showing three main sections of settings:

- Resiliency Manager Settings:**
 - Resiliency Manager Instance Name: vrp-8M
 - Resiliency Manager Hostname: vrp-8m
 - Resiliency Manager Instance Size: 1x Standard D8s v3 (8 vcpus, 32 GB memory)
 - Data Disk Size in GB: 100
 - Data disk type: Premium SSD
- Infrastructure Management Server (IMS) Settings:**
 - IMS Instance Name: vrp-ims
 - IMS Hostname: vrp-ims
 - IMS Instance Size: 1x Standard F8s (8 vcpus, 16 GB memory)
 - Data Disk Size in GB: 40
 - Data disk type: Premium SSD
- Replication Gateway Settings:**
 - Replication Gateway Instance Name: vrp-gw1
 - Replication Gateway Hostname: vrp-gw1
 - Replication Gateway Instance Size: 1x Standard F8s (8 vcpus, 16 GB memory)
 - Staging Disk Size in GB: 50
 - Staging disk type: Premium SSD
 - Enable Ultra disk compatibility:
 - Availability Zone: (empty)

Table 2-18 Advanced settings for Azure deployment

Input field	Description
Resiliency Manager Settings	
Instance Name	Name of the appliance instance
Host name	Hostname of the appliance. This is changed and allocated from DHCP details
Instance Size	Size of the appliance instance
Data Disk Size in GB	Additional disk of minimum 100 GB required for Resiliency Manager
Infrastructure Management Server (IMS) Settings	
Instance Name	Name of the appliance instance

Table 2-18 Advanced settings for Azure deployment (*continued*)

Input field	Description
Host name	Hostname of the appliance. This is changed and allocated from DHCP details
Instance Size	Size of the appliance instance
Data Disk Size in GB	Additional disk of minimum 40 GB required for IMS
Replication gateway Settings	
Instance Name	Name of the appliance instance
Host name	Hostname of the appliance. This is changed and allocated from DHCP details
Instance Size	Size of the appliance instance
Staging disk size in GB	Additional disk of minimum 50 GB required for Replication Gateway
Staging disk type	Choose staging disk type for Replication Gateway. Learn more about the managed disk types. Managed disk types
Enable Ultra disk compatibility	Choose this check box if you want to support ultra disks on your Replication Gateway instance
Availability Zone	Specify an availability zone if the selected region for deploying the Replication Gateway has availability zones. Keep this field empty if the region does not support availability zones. The deployment can fail if an incorrect input is provided. See Using Azure ultra disks

Figure 2-6 Network settings for Azure deployment

Basics
 Advanced
 Resiliency Manager Network
 Infrastructure Management Network

Is Multi-NIC Deployment? Yes
 No

i Select virtual network where you want to deploy Resiliency Manager. We do not support creation of virtual network as part of deployment. Make sure you select one of the existing networks.

Configure virtual networks

Network Name * (new) VRPvNet
[Create new](#)

Subnet * (new) default (10.2.0.0/24)

Is the NIC eth0 behind NAT? Yes
 No

Table 2-19 Network settings for Azure deployment

Input field	Description
Is Multi-NIC Deployment?	You can deploy Resiliency Manager, IMS, or Replication Gateway with multiple network interfaces. These interfaces can be used for different communication purpose. Select Yes if required..
Are Both NICs in Same Subnet	Select Yes if multiple interfaces are part of same subnet of a virtual network.
Network Name	Specify the name of an existing network.
Select Subnet	Select a subnet to be associated with the virtual appliance.

Table 2-19 Network settings for Azure deployment (*continued*)

Input field	Description
Role of NICs'	In case of multiple interfaces, select NIC for its communication purpose. Make sure you select 'eth0' for at least one intent. If it is single NIC deployment then same NIC is used.
Is the NIC eth0 behind NAT?	Select Yes if NAT is used in your network setup and provide NAT IP and fully qualified NAT hostname.

Similar input is required for configuring IMS, Replication Gateway.

Deploy virtual appliances in Azure Stack using Azure Stack Marketplace

To know about virtual appliance deployment in Veritas Resiliency Platform

A few constraints are applied when you deploy Resiliency Platform in Azure Stack:

See [“Constraints when you deploy the virtual appliances in Azure Stack”](#) on page 405.

Note that, Azure Stack is supported with Resiliency Platform version 3.5 and above.

Below are few prerequisites before deploying in Azure Stack Marketplace:

- Your Azure Stack deployment must have internet connectivity.
- Azure Stack needs to be registered with Azure.

To download Azure Marketplace items to Azure Stack:

- 1 Sign in to the Azure Stack administrator portal.
- 2 Review the available storage space before downloading marketplace items. Later, when you select items for download, you can compare the download size to your available storage capacity.

To review available space, navigate to **Region management** and select the region you want to explore. Later, go to **Resource Providers > Storage**.
- 3 Open Azure Stack Marketplace and connect to Azure. To do so, navigate to: **Marketplace management service** and select Marketplace items, and then select **Add from Azure**.
- 4 Select the **Veritas™ Resiliency Platform Express Install** from the list and then select **Download**. The download time may vary depending upon the internet connectivity. After the download completes, you can deploy the new marketplace item either as an Azure Stack operator or a user.

After the download is complete, to deploy the virtual appliances, refer to the below topic:

See [“Deploying the virtual appliances in Azure through Azure Marketplace”](#) on page 405.

Deploying the virtual appliances in vCloud

To know about virtual appliance deployment in Veritas Resiliency Platform:

You need to deploy at least three Resiliency Platform virtual appliances in your data center in vCloud and then configure Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway.

To deploy the Resiliency Platform virtual appliances in vCloud

1 Recommendations:

While deploying the virtual appliances, ensure the following:

- Deploy the Resiliency Manager and IMS together in a vApp. This vApp will function as a Main Management vApp.
- Deploy the Replication Gateway in a vApp other than the one in which you have deployed the Resiliency Manager and IMS.
- Deploy the Replication Gateway in the Organization where you want to migrate the virtual machines from your on-premises data center.

2 Open vCloud Director, log into the Organization where you want to deploy the Resiliency Manager, IMS, or Replication Gateway.

3 Click **My Cloud**, in the left pane click **vApps**.

4 Use one of the following methods to deploy the OVA in vCloud:

- **Add vApp From Catalog:** Select the OVA from the catalog while adding the vApp.
- **Build a new vApp:** This method can be used to deploy multiple virtual machines at a time. Select the OVA from the catalog while building a new vApp.

Follow the documentation of VMware to deploy the OVA in vCloud using any one of the above-mentioned methods.

Ensure that the details on **Customize Hardware** page match the system resource requirements mentioned:

- 5 After the successful deployment of the OVA, configure the appliance as a Resiliency Manager, IMS, or Replication Gateway.

See “[Configuring the Resiliency Manager or IMS](#)” on page 439.

See “[Configuring the Replication Gateways](#)” on page 445.

Prerequisites for deploying the virtual appliances in vCloud Director

Following are the deployment prerequisites for recovery to vCloud:

- Configure vCloud to have unique BIOS UUID for all the virtual machines when instantiating from a vApp template. By default, all the virtual machines that are created when you deploy a vApp template are assigned the same BIOS UUID. To change this default behavior, follow the steps given in the VMware knowledge bank article:
<https://kb.vmware.com/kb/2002506>
- Set the value of the **disk.enableUUID** configuration parameter as **True** for the following templates uploaded to catalog:
 - Veritas_Resiliency_Platform_VMware_vCloud_vApp_Template_Site1
 - Veritas_Resiliency_Platform_VMware_vCloud_vApp_Template_Site2
 - Veritas_Resiliency_Platform_VMware_vCloud_vApp_Template_Win_Site1
 - Veritas_Resiliency_Platform_VMware_vCloud_vApp_Template_Win_Site2
 - Resiliency Platform Data Mover
- If you want to get the virtual appliances deployed on some specific datastores in vCloud, then you need to create a storage policy. If you do not create a storage policy, the virtual appliances are deployed on any of the available datastores.

Uploading OVA files to catalog

The cloud administrator needs to create a catalog and upload the OVA files in to the catalog.

To upload the OVA files to catalog

- 1 Log into the vCloud service provider Organization in vCloud Director.
- 2 Create a catalog and share the catalog with the tenant Organizations.
- 3 Upload the OVA files of all the Resiliency Platform appliances and all the vApp templates in to the catalog that you have created.
- 4 Ensure that the template vApp names are set to the following:

- VRP_VAPP_TEMPLATE_SITE1
 - VRP_VAPP_TEMPLATE_SITE2
 - VRP_VAPP_TEMPLATE_WIN_SITE1
 - VRP_VAPP_TEMPLATE_WIN_SITE2
- 5 Make the OVAs available to the tenants for deployment of the virtual appliances.
 - 6 Edit the settings for Data Mover and VRP_VAPP_TEMPLATE virtual machines from underlying vCenter servers to set the **disk.enableUUID** parameter as **True**.

Deploying the virtual appliances in Orange Recovery Engine

To know about virtual appliance deployment in Veritas Resiliency Platform:

Following is an overview of the key steps that are performed for deploying the Resiliency Platform virtual appliances in Orange Recovery Engine. You need to perform these steps for each of the virtual appliances:

Table 2-20 Overview of deployment process in Orange Recovery Engine

Step	Action	Description
1	Ensure that the prerequisites for deploying the virtual appliances in Orange Recovery Engine are met.	See “Prerequisites for deploying the virtual appliances in Orange Recovery Engine” on page 416.
2	Download the zip files required for deploying the Resiliency Platform virtual appliance in Orange Recovery Engine.	
3	Upload the image files to Object Storage Service (OBS) using OBS Browser tool. Alternatively, you can use any S3 compliant tool to perform this task.	See “Uploading image files to OBS” on page 416.
4	Using the Orange Recovery Engine console, register the system disk image as a private image.	See “Registering the disk image as a private image” on page 417.
5	Using the Orange Recovery Engine console, launch the ECS instance and attach the data disk to the system disk image that you registered in the previous step.	See “Launching the system disk image instance” on page 418.
6	Restart the ECS.	

Prerequisites for deploying the virtual appliances in Orange Recovery Engine

Following are the prerequisites for deploying the Resiliency Platform virtual appliances in Orange Recovery Engine:

- Ensure to create virtual private network (VPC) and subnets in Orange Recovery Engine.
- Ensure to configure an external DNS server in Orange Recovery Engine. This DNS server is used while configuring the virtual appliances. The host names of IMS and Replication Gateway at source data center should be resolvable from the target data center, and hostname of Resiliency Manager should be resolvable from the source data center.
- Follow the documentation of Orange Recovery Engine to create the required security groups. Make sure that the security groups meet the network and port requirements mentioned in the Resiliency Platform documentation are open.
- Ensure that Resiliency Manager and Infrastructure Management Server (IMS) have outgoing internet access enabled. You may choose to restrict incoming internet access on these virtual appliances.
- OBS Browser tool needs to be installed on the system through which you are planning to deploy the virtual appliances in Orange Recovery Engine. Alternatively, you can use any S3 compliant tool to perform this task.
- The user who is going to deploy the virtual appliances must have read and write access control list (ACL) on the OBS bucket that is used for deploying the virtual appliances.
- Ensure that the bucket policy settings for the bucket that is used for deploying the virtual appliances do not restrict uploading files to the bucket.

See [“Deploying the virtual appliances in Orange Recovery Engine”](#) on page 415.

Uploading image files to OBS

Before you start deploying the Resiliency Platform virtual appliances in Orange Recovery Engine, you need to upload the downloaded image files to Object Storage Service (OBS).

To upload image files to OBS

- 1 Log in to OBS Browser.
- 2 Select a bucket with **Standard** storage class. If such a bucket does not exist, create a bucket.
- 3 Select the bucket where you want to upload the files.

- 4 Click **Upload** and then click **Upload File**.
- 5 In the next wizard, select the image files to be uploaded.
- 6 Ensure that **Standard** storage class is selected.
- 7 Click **OK** to upload the files. If an upload operation gets suspended or fails, restart the operation. The task will be resumed from the point where it got suspended last time.

See [“Deploying the virtual appliances in Orange Recovery Engine”](#) on page 415.

Registering the disk image as a private image

As a part of deploying the Resiliency Platform virtual appliances in Orange Recovery Engine, you need to register the system image file as a private image.

To register the system image file as a private image

- 1 Log in to the Orange Flexible Engine management console and go to **Image Management Service**.
- 2 On the **Image Management Service** page, click **Create Private Images** tab and click **Create Image**.
- 3 On the **Create Image** page, under **Image Type and Source**, select **System Disk Image** as type.
- 4 For **Source**, select **Image file** to use an external image file.
- 5 Select the bucket where you have uploaded the image files. Navigate to the system image file that you want to register as a private image and select it:
 - System disk for Resiliency Manager:
`Resiliency_Platform_RM_v36-disk1.qcow2`
 - System disk for IMS: `Resiliency_Platform_v36-disk.qcow2`
 - System disk for Replication Gateway:
`Resiliency_Platform_DM_v36-disk1.qcow2`
- 6 Under **Image Information**, ensure that **ECS system disk image** is selected as **Function**.
- 7 For **System Disk (GB)**, enter a minimum value of 40 GB.
- 8 Enter a name for the image.
- 9 Click **Create Now** and review the information displayed on the next page.
- 10 Click **Submit** button to agree to the Orange Flexible Engine Image Disclaimer and click **Submit**.

See [“Deploying the virtual appliances in Orange Recovery Engine”](#) on page 415.

Launching the system disk image instance

After registering the system disk as a private image in Orange Recovery Engine, you need to launch the system disk image instance and attach the data disk to the system disk. For Resiliency Manager and IMS, you need to attach the data disk image that you had downloaded earlier. For Replication Gateway, you attach an extra disk of minimum 50 GB.

After registering the system disk as a private image in Orange Recovery Engine, you need to launch the system disk image instance and attach the data disk to the system disk. For Resiliency Manager, IMS and Replication Gateway attach an extra external disk with below minimum size:

Resiliency Manager: 100 GB

IMS: 40 GB

Replication Gateway: 50 GB

To launch the system disk image instance

- 1 In the Orange Recovery Engine management console, go to **Image Management Service**.
- 2 Select the private image that you had registered in the earlier step.
- 3 Click **Apply for Server** displayed next to the name of the private image.
- 4 On the next page, select the billing mode that suits your requirements.
- 5 Ensure that the desired region and Availability Zone are selected.
- 6 Specify the type and select a flavor that matches the system resource requirements for various Resiliency Platform components.

Note: Select a flavor which is KVM hypervisor based for Replication Gateway, as based hypervisor is not supported.

For more information on hypervisor flavor,

see https://docs.prod.cloud-ocb.orange-business.com/en-us/usermanual/ec/en-us_topic_0035470101.html

- 7 In the **Disk** section, select **High I/O** or **Ultra-high I/O** for the system disk. It is strongly recommended to select **Ultra-high I/O** for the disk.
- 8 Ensure that the required image name is selected in the drop down list for image name.
- 9 Click **Add data Disk** and then select **High I/O** or **Ultra-high I/O** for the data disk.

- 10** You can attach external disk of below minimum size to the virtual appliances:
Resiliency Manager: 100 GB
IMS: 40 GB
Replication Gateway: 50 GB
 - 11** Select the **VPC** that you want to use for the deployment.
 - 12** Enter the value for **Security Group** and **NIC**. It is strongly recommended not to assign any public IP or **EIP** to any of the Resiliency Platform appliances.
 - 13** You need to create a key pair or select existing key pair if you have already created it.
 - 14** Enter a name for the ECS. Click **Next**.
 - 15** Verify the details displayed under **Configuration** and click **Create Now**.
- See [“Deploying the virtual appliances in Orange Recovery Engine”](#) on page 415.

Deploying the virtual appliance through VMware vSphere Client

To know about virtual appliance deployment in Veritas Resiliency Platform:

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

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

11 Power on the virtual machine.

To deploy Resiliency Platform through VMware vSphere Web Client

- 1** Download the VMware supported OVA file for the Resiliency Platform virtual appliance on a system where VMware vSphere Web Client is installed.
- 2** In the VMware vSphere Web Client, click **vCenter Servers** and select a vCenter Server. Click **Actions > Deploy OVF template**.
- 3** Select the source location of the Resiliency Platform virtual appliance OVA file.
- 4** Specify a name and location for the deployed template.
- 5** Select a cluster, host, vApp, or resource pool in which to run the deployed template.
- 6** Select a location to store the files for the deployed template.
- 7** Configure the networks the deployed template should use.
- 8** Review the virtual machine configuration and click **Finish**.
- 9** If you want to use DHCP as your network, enter the MAC address of the appliance in the DHCP server. For information on obtaining the MAC address of the appliance, see the documentation of VMware vSphere client.
- 10** Power on the virtual machine.

You can now configure the Resiliency Platform component.

Note: When you deploy the Resiliency Platform Data Mover appliance from vCenter Server 6.5, a warning related to advanced configuration may be displayed. You can ignore the warning and click next to accept the advanced configuration options.

For more information on VMware vSphere Desktop Client or VMware vSphere Web Client, refer to VMware documentation.

See [“About configuring the Resiliency Platform components”](#) on page 437.

Deploying the virtual appliance through Hyper-V Manager

To know about virtual appliance deployment in Veritas Resiliency Platform:

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.
- 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 for IMS or Replication Gateway and 32 GB RAM for Resiliency Manager.
- 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 the VHD file of the Resiliency Platform virtual appliance as **IDE Controller 0**.
- 11 Click **Apply**, and then click **OK**.
- 12 If you want to use DHCP as your network, enter the MAC address of the appliance in the DHCP server. For information on obtaining the MAC address of the appliance, see the documentation of Hyper-V Manager.
- 13 Right-click the name of the virtual machine and select **Start** to power on the virtual machine.

You can now configure the Resiliency Platform component.

See [“About configuring the Resiliency Platform components”](#) on page 437.

Deploying the virtual appliances in Google Cloud Platform (GCP) through GCP Marketplace

Veritas Resiliency Platform enables you to deploy the virtual appliances in Google Cloud Platform through GCP Marketplace. There are four offerings available for deploying the virtual appliances using the templates:

- **Veritas™ Resiliency Platform Express Install 10.0:** Installs a Resiliency Manager, an IMS, and a Replication Gateway appliance of version 10.0.
- **Veritas™ Resiliency Manager 10.0:** Installs a Resiliency Manager appliance of version 10.0.

- **Veritas™ Infrastructure Management Server 10.0:** Installs an Infrastructure Management Server appliance of version 10.0.
- **Veritas™ Replication Gateway 10.0:** Installs a Replication Gateway appliance of version 10.0.

To deploy the virtual appliances in GCP using the templates

1 Prerequisites:

Ensure that the prerequisites are met. Refer to [API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace](#), for deploying the Resiliency Manager and IMS in Google Cloud Platform Marketplace.

Refer See “[Prerequisites for deploying the virtual appliances in Google Cloud Platform](#)” on page 429.

Refer See “[Ports required for recovery of assets to Google Cloud Platform](#)” on page 430.

- 2 Go to the **GCP Marketplace** and search for the offerings.
- 3 Select the desired offerings from the four available offers. GCP marketplace lets you launch the selected offering. Click **Next**.
- 4 On the next page, provide the values for the input fields: See “[Providing inputs for Resiliency Platform template](#)” on page 423.
- 5 Click **Deploy** to start the deployment of resources. Once the deployment starts you will see list of resources being creating as part of deployment.

Next Steps:

SSH to virtual appliance to set admin password:

Once the deployment succeeds you need to SSH to all the servers and change the default password. Resiliency Platform virtual appliances comes with default password for ‘admin’ user and expires at first login. Follow the below steps:

1. Select the virtual machine and enable port 22 in firewall policy if not selected as part of deployment.
2. Login to the virtual machine using ‘Connect to serial console’ with username ‘admin’ and with default password or SSH public key if you have provided during deployment.

Access Resiliency Manager Web UI:

Once you have set the password you can access the Resiliency Manager Web UI using URL: https://<Resiliency_manager_hostname>/ and perform the operations.

Providing inputs for Resiliency Platform template

You need to provide inputs for creating instances using templates. Some of the fields get auto populated with the default value, you can change the values if required. For rest of the parameters, you need to enter a valid value.

Table 2-21 Input required for creating configuration for Resiliency Manager

Field	Description
Instance Count	Number of Resiliency Managers to be created. Only one instance per deployment is currently supported.
Machine type and Series	Select configuration of the instance. Minimum 8 vCPU and 32 GB RAM is required.
Boot disk size in GB	Shows the OS disk size. For the instance 50 GB is required
Boot Disk type	Select disk type for OS disk
Data disk type	Select disk type of data disk. The instance comes with one data disk.
Data disk size in GB	Disk size for data disk. The default value is 100, you can increase the size as required.
Network interfaces	Select Network configuration for NICs. Resiliency Manager allows to have max 2 NICs
Allow TCP port 22 from the Internet	Select checkbox if you want to enable SSH port on the VM, also provide CIDR block which needs to be allowed by default. If kept empty, then SSH will be allowed to all, i.e. 0.0.0.0/0
Instance Name	Name of the instance to be used. If DNS is not configured, then this will be set as short hostname.
Create role with ID VRP_ROLE_RM	Provide custom role with permissions.
Create service account	Provide service account for IAM policy with the role.
Service Account name	Provide service account name having following permissions required for Resiliency Manager to work: <ul style="list-style-type: none">■ compute.regions.list■ storage.buckets.list■ runtimeconfig.configs.create■ runtimeconfig.variables.create■ runtimeconfig.waiters.create

Table 2-21 Input required for creating configuration for Resiliency Manager
(continued)

Field	Description
NIC to be used for communication with other Resiliency Managers	In case of multiple NICs communication with other Resiliency Manager can be restricted on a particular NIC. Select an Interface to be used for the communication.
NIC to be used for communication with IMS	In case of multiple NICs communication with IMS can be restricted on a particular NIC. Select an Interface to be used for the communication.
NIC to be used for accessing the User Interface	In case of multiple NICs Web UI can be restricted on a particular or allowed on all NIC. Select an Interface to be used for the communication.
NIC to be used as default gateway	Select the interface to be used as default gateway for external communication
Is the eth0 Network Interface behind NAT?	If NAT is configured for eth0 interface, then select 'True'. You need to provide NAT hostname and IP address.
Resiliency Manager eth0 NAT Hostname (Optional)	Provide NAT hostname for eth0 if NAT is configured.
Resiliency Manager eth0 NAT IP (Optional)	Provide NAT IP for eth0 if NAT is configured.
Is the eth1 Network Interface behind NAT?	If NAT is configured for eth1 interface, then select 'True'. You need to provide NAT hostname and IP address.
Resiliency Manager eth1 NAT Hostname (Optional)	Provide NAT hostname for eth1 if NAT is configured.
Resiliency Manager eth1 NAT IP (Optional)	Provide NAT IP for eth1 if NAT is configured.

Table 2-22 Input required for creating configuration for Infrastructure Management Server (IMS)

Field	Description
Instance Count	Number of IMS to be created. Only one instance per deployment is currently supported.
Machine type and Series	Select configuration of the instance. Minimum 8 vCPU and 32 GB RAM is required.
Boot disk size in GB	Shows the OS disk size. For the instance 30 GB is required
Boot Disk type	Select disk type for OS disk

Table 2-22 Input required for creating configuration for Infrastructure Management Server (IMS) (*continued*)

Field	Description
Data disk type	Select disk type of data disk. The instance comes with one data disk.
Data disk size in GB	Disk size for data disk. Default is 40 but you can increase as required.
Network interfaces	Select Network configuration for NICs. Resiliency Manager allows to have max 2 NICs
Allow TCP port 22	Select checkbox if you want to enable SSH port on the VM, also provide CIDR block which needs to be allowed by default. If kept empty then SSH will be allowed to all, i.e. 0.0.0.0/0
Instance Name	Name of the instance to be used. If DNS is not configured, then this will be set as short hostname.
Create role with ID VRP_ROLE_IMS	Provide custom role with permissions.
Create service account	Provide service account for IAM policy with the role.
Service Account name	Provide service account name having following permissions required for operations. Check the prerequisites list for required permissions.
NIC to be used for communication with Resiliency Managers	In case of multiple NICs communication with other Resiliency Manager can be restricted on a particular NIC. Select an Interface to be used for the communication.
NIC to be used for communication with Replication Gateway	In case of multiple NICs communication with Replication Gateway can be restricted on a particular NIC. Select an Interface to be used for the communication.
Network Interface to be used as the default gateway	Select the interface to be used as default gateway for external communication
Is selected NIC to communicate Resiliency Manager behind NAT?	NAT can be configured on the NIC which is used for communication with Resiliency Manager. Select True if NAT is configured on the NIC and provide NAT hostname and IP
NAT Hostname (Optional)	Provide NAT hostname if NAT is configured.
NAT IP (Optional)	Provide NAT IP if NAT is configured.

Table 2-23 Input required for creating configuration for Replication Gateway

Field	Description
Instance Count	Number of IMS to be created. Only one instance per deployment is currently supported.
Machine type and Series	Select configuration of the instance. Minimum 8 vCPU and 32 GB RAM is required.
Boot disk size in GB	Shows the OS disk size. For the instance 30 GB is required
Boot Disk type	Select disk type for OS disk
Data disk type	Select disk type of data disk. The instance comes with one data disk.
Data disk size in GB	Disk size for data disk. Default is 50 but you can increase as required.
Network interfaces	Select Network configuration for NICs. Resiliency Manager allows to have max 2 NICs
Allow TCP port 22	Select checkbox if you want to enable SSH port on the VM, also provide CIDR block which needs to be allowed by default. If kept empty then SSH will be allowed to all, i.e. 0.0.0.0/0
Instance Name	Name of the instance to be used. If DNS is not configured, then this will be set as short hostname.
Create role with ID VRP_ROLE_GW	Provide custom role with permissions.
Create service account	Provide service account for IAM policy with the role.
Service Account name	Provide service account name having following permissions required for operations. Check the prerequisites list for required permissions.
NIC to be used for communication with peer Replication Gateway	In case of multiple NICs communication with other Resiliency Manager can be restricted on a particular NIC. Select an interface to be used for the communication.
NIC to be used for communication with IMS	In case of multiple NICs communication with Infrastructure Management Server can be restricted on a particular NIC. Select an Interface to be used for the communication.
NIC to be used for communication with workload Virtual Machines	In case of multiple NICs communication with workload can be restricted on a particular NIC. Select an interface to be used for the communication.

Table 2-23 Input required for creating configuration for Replication Gateway
(continued)

Field	Description
NIC to be used as the default gateway	Select the interface to be used as default gateway for external communication
NIC to communicate with peer Replication Gateway behind NAT?	NAT can be configured on the NIC which is used for communication with peer Replication Gateway. Select True if NAT is configured on the NIC and provide NAT hostname and IP address.
NAT Hostname (Optional)	Provide NAT hostname if NAT is configured.
NAT IP (Optional)	Provide NAT IP address if NAT is configured.

Table 2-24 Common inputs

Fields	Inputs required
SSH public key to this instance	Provide public SSH key to allow key based authentication to the instances. The same key will be configured on all the instances
NTP Server	Provide comma separated one or more NTP servers' hostname or IP address. All the Resiliency Platform virtual appliances must in time sync to perform operations.
Timezone	Select timezone for the instances.

API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace.

Following are the API permissions required to deploy Resiliency Manager and IMS using Google Cloud Platform through Marketplace. Before referring to the below table, you need to know there are some more permissions. Refer to the topic [API permissions required for Google Cloud Platform](#)

Table 2-25 Permissions for deploying Resiliency Manager and IMS using Google Cloud Platform through Marketplace.

Service name	Permissions
runtimeconfig	runtimeconfig.configs.create
	runtimeconfig.variables.create
	runtimeconfig.waiters.create

Prerequisites for deploying the virtual appliances in Google Cloud Platform

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

1. Follow the documentation of Google Cloud Platform to create the required network tags. Make sure that the network tags meet the network and port requirements mentioned in the Resiliency Platform documentation are open for communication. If you deploy Resiliency Platform components through Google Cloud Platform marketplace, the required network tags are automatically created.

See [“Ports required for recovery of assets to Google Cloud Platform”](#) on page 430.
2. Create individual Service Accounts for Resiliency Manager and IMS with certain permissions. These service accounts are used for authenticating the operations performed by the Resiliency Platform components in Google Cloud Platform.

See [“API permissions required for Google Cloud Platform”](#) on page 460.
3. If you deploy Resiliency Platform through Google Cloud Platform marketplace, then the required service accounts are automatically created through Google Cloud Platform template used by marketplace deployment.
4. Ensure that there is direct communication between the premise network and the Google Cloud Platform network. It is recommended to use VPN for Google Cloud Platform environment.
5. Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.

Deploying the virtual appliances in Google Cloud Platform using OVA files

This topic explains about the key steps that are performed for deploying the Resiliency Platform virtual appliances in Google Cloud Platform (GCP). To know about virtual appliance deployment in Veritas Resiliency Platform, refer

Table 2-26 Overview of deployment process in GCP

Step	Action	Description
1	Ensure that the prerequisites for deploying virtual appliances in GCP are met.	Refer See " Prerequisites for deploying the virtual appliances in Google Cloud Platform " on page 429.
2	Upload the OVA files to Google Cloud Storage.	Refer See " Uploading the OVA file using web-based method " on page 433. Refer See " Uploading the OVA file using command-line method " on page 434.
3	Create image from the uploaded OVA file.	Refer See " Creating Image using web-based method " on page 435. Refer See " Creating Image using command-based method " on page 435.
4	Launch the instances of virtual appliances to deploy Resiliency Manager, Infrastructure Manager (IMS), and Replication Gateway	Refer See " Launching the instances of virtual appliances " on page 436.

Prerequisites for deploying the virtual appliances in Google Cloud Platform

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

1. Follow the documentation of Google Cloud Platform to create the required network tags. Make sure that the network tags meet the network and port requirements mentioned in the Resiliency Platform documentation are open for communication. If you deploy Resiliency Platform components through Google Cloud Platform marketplace, the required network tags are automatically created.

See "[Ports required for recovery of assets to Google Cloud Platform](#)" on page 430.

2. Create individual Service Accounts for Resiliency Manager and IMS with certain permissions. These service accounts are used for authenticating the operations performed by the Resiliency Platform components in Google Cloud Platform.

- See [“API permissions required for Google Cloud Platform”](#) on page 460.
3. If you deploy Resiliency Platform through Google Cloud Platform marketplace, then the required service accounts are automatically created through Google Cloud Platform template used by marketplace deployment.
 4. Ensure that there is direct communication between the premise network and the Google Cloud Platform network. It is recommended to use VPN for Google Cloud Platform environment.
 5. Ensure to deploy the IMS in the region to which you plan to associate the cloud data center.

Ports required for recovery of assets to Google Cloud Platform

Following is the list of ports required for recovery of assets to Google Cloud Platform:

Table 2-27 Ports required for recovery of assets to Google Cloud Platform

Ports for recovery to Google Cloud Platform
See “Ports required for Resiliency Manager” on page 78.
See “Ports required for IMS” on page 80.
See “Ports required for Replication Gateway used for recovery to cloud data center” on page 81.
See “Ports required for hosts” on page 82.

API permissions required for Google Cloud Platform

Following are the permissions required for the roles that you need to create for Resiliency Manager and IMS for recovery to Google Cloud Platform data center. If you are deploying your appliances through Marketplace, then refer to

[API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace.](#)

Table 2-28 API Permissions required for role for Resiliency Manager

Service name	Permissions
compute	compute.regions.list
storage	storage.buckets.list

Table 2-29 API Permissions required for role for IMS

Service name	Permission
cloudkms	cloudkms.cryptoKeyVersions.list
	cloudkms.cryptoKeys.list
	cloudkms.keyRings.list

Table 2-29 API Permissions required for role for IMS (*continued*)

Service name	Permission
compute	compute.addresses.list
	compute.diskTypes.list
	compute.disks.create
	compute.disks.createSnapshot
	compute.disks.delete
	compute.disks.list
	compute.disks.use
	compute.disks.get
	compute.firewalls.list
	compute.globalOperations.list
	compute.images.create
	compute.images.get
	compute.images.useReadOnly
	compute.instances.attachDisk
	compute.instances.create
	compute.instances.delete
	compute.instances.detachDisk
	compute.instances.get
	compute.instances.list
	compute.instances.setMetadata
	compute.instances.setTags
	compute.instances.start
	compute.instances.stop
compute.machineTypes.list	
compute.networks.list	

Table 2-29 API Permissions required for role for IMS (*continued*)

Service name	Permission
	compute.projects.get
	compute.regionOperations.list
	compute.snapshots.create
	compute.snapshots.delete
	compute.snapshots.list
	compute.snapshots.useReadOnly
	compute.subnetworks.list
	compute.subnetworks.use
	compute.zoneOperations.list
	compute.zones.list
	compute.addresses.useInternal
storage	storage.objects.create
	storage.objects.get

Table 2-30 Permissions required to discover Shared VPC for a role for IMS on the host project.

Service name	Permission
compute	compute.networks.list
	compute.firewalls.list
	compute.addresses.list

Note: To share subnets from the host project with configured project, you need to add the Service Account associated with IMS as **Principal** and provide a **Compute Network User role** on the shared subnets.

Uploading the OVA file using web-based method

You can create a Google Cloud Platform cloud storage bucket and upload the ova file to that bucket using a web-based method.

To upload the OVA file using web-based method

- 1 Log into the Google Cloud Platform console.
- 2 Navigate to **Cloud Storage** and click **Create Bucket**.
- 3 Enter a name for the bucket and select the appropriate region.
- 4 Select the Standard storage class.
- 5 Click **Create**.
- 6 Once the bucket is created, open the bucket and click **Upload Files**.
- 7 Select the OVA files from your local disk and it will start uploading the file.

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Uploading the OVA file using command-line method

You need to first create a Cloud Storage bucket in Google Cloud Platform and then upload your ova file to that bucket.

To upload the OVA file using command-line method

- 1 Download and install the [GCP Command Line Interface](#).
- 2 Use the `gsutil mb` command to create a new bucket. Bucket names must be unique and should be DNS compliant:

```
gsutil mb gs://<bucket_name>
```

Where:

- `bucket_name` is the name you want to give your bucket.
- If the request is successful, the command returns the following message:

```
Creating gs://<bucket_name>/...
```

- 3 Upload the OVA file by executing the following command:

```
gsutil cp <object_location> gs://<destination_bucket_name>/
```

Where:

- `object_location` is the local path to your OVA file. For example, `Desktop/RM.ova`.
- `destination_bucket_name` is the name of the bucket to which you are uploading your object. For example, `my-bucket`.

If the request is successful, the response looks like the following example:

```
Operation completed over 1 objects/58.8 KiB
```

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Creating Image using web-based method

Once you upload the OVA files to Google Cloud Storage bucket, you need to create a Image from the OVA files that you have uploaded. This image can be later used to launch the instances for deploying Resiliency Manger, Infrastructure Manager in Google Cloud Platform.

To create image

- 1 Login to the Google Cloud Platform console.
- 2 Navigate to **Images** and click on **Create Image**.
- 3 Enter a name for image.
- 4 Select the **Source** as **Virtual Disk**.
- 5 In Cloud Storage file, browse the OVA uploaded into your bucket.
- 6 In Operating System on virtual disk, select **No operating system. Data only**.
- 7 Click **Create**.

It will start importing image from the selected OVA file.

Note: If your default VPC setting for “Subnet creation mode” from the project is set to “Custom subnets” then create image will not work from Google Cloud Platform console. In that case, you need to use command-based method only.

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Creating Image using command-based method

You can also create image using command-based method:

To create an image by command line

- 1 Open command prompt.
- 2 Execute the following command to create the image.

```
gcloud compute images import <image_name> --source-file
gs://<bucket_name>/<file_name> --no-guest-environment --no-address
--network projects/<project_name>/global/networks/default --subnet
projects/<project_name>/regions/<region>/subnetworks/default
```

Where:

- `bucket_name` is the name of the bucket in which you have uploaded the OVA file.
- `project_name` is name of the project for which you want to create the image.
- `Region` is name of the region for which you need to create the Resiliency Platform virtual appliances.
- `File name` is the name of the ova file which is uploaded.

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

Launching the instances of virtual appliances

Once an image gets created, you can use the image to launch instances to deploy the Resiliency Manager and any number of Infrastructure Management Servers (IMS), and Replication Gateways in Google Cloud Platform.

To launch the instances of virtual appliances

- 1 Go to the Google Cloud Platform console and navigate to **Images** under **Storage**.
- 2 Click on the image which you want to select and click **Create Instance**.
- 3 Make sure to select appropriate Machine Type that matches with the system resource requirements mentioned in the documentation.

Network Optimization should be high for the instance.
- 4 Select the required Region and Zone in which you wish to deploy the Resiliency Platform.
- 5 Under **Identity and API access**, select a Service Account created with all the necessary permissions for Resiliency Manager and IMS.

If you deploy the appliances using marketplace, then template deployment will create the required service account and associate it with the instance.
- 6 Under **Networking**, enter the Network Tag created for the particular instance.
- 7 Select the proper Network and Subnetwork in which you want to create the setup.
- 8 Under Disks, you need to attach an extra disk of size 50 GB for Resiliency Manager, 40 GB for IMS and 50 GB for Replication Gateway.
- 9 Click **Create** to launch the instance.

In the **Select an existing key pair** or **Create a new key pair wizard**, you can choose an existing key pair, or create a new one. If you create a new key pair, ensure to click the **Download key pair** button and download. You will need the

private key from this key pair to login as admin user for completing the bootstrap process.

See [“Deploying the virtual appliances in Google Cloud Platform using OVA files”](#) on page 429.

About configuring the Resiliency Platform components

After the Veritas Resiliency Platform virtual appliance deployment, you are expected to configure the Resiliency Platform component that you have deployed, through the bootstrap process. The bootstrap process is automatically invoked when you log in to the virtual appliance console for the first time using the admin user login.

Note: There is no sequence required for configuring the Resiliency Platform components. You can configure the components in any sequence on source as well as target data centers. Only after configuring Resiliency Manager, a URL for the Resiliency Platform web console login is provided and then you can access that URL in a web browser to log in to the web console.

The following settings are configured as part of this process to set up the component:

- **Host Network settings:** Settings such as fully qualified hostname (FQDN), IP address, subnet mask, default gateway, and DNS server. Before you use the hostname and the IP address, you need to register them with the DNS server.
- **Appliance settings:** Settings such as NTP server.
- **Product settings:** Configures the virtual appliance as a Resiliency Manager, Infrastructure Management Server (IMS), Replication Gateway.

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

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

Before configuring the component through product bootstrap, ensure that the prerequisites are met.

See [“Prerequisites for configuring Resiliency Platform components”](#) on page 438.

See [“Configuring the Resiliency Manager or IMS”](#) on page 439.

See [“Configuring the Replication Gateways”](#) on page 445.

Prerequisites for configuring Resiliency Platform components

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

- Veritas Resiliency Platform now supports Internet protocol version 6 (IPv6) along with Internet protocol version 4 (IPv4) .
- Before you use the hostname and the IP address in the **Network settings**, you need to register them with the DNS server. Also ensure that the reverse lookup for that IP address works.
- Ensure that the host name corresponding to the IP address is less than 64 characters long. In case of NAT, this is required for host names corresponding to both private and public IP addresses.
- If you plan to use the DHCP server, the DHCP server should be reachable from the subnet and should be able to respond to the subnet where you plan to deploy the product. This requirement is also applicable to static DHCP.
- To use DHCP network, you need to reserve an IPv4 address for the virtual appliance in the DHCP server along with the corresponding MAC address. You cannot configure IPv6 address for the virtual appliance in the DHCP server.
- If you are configuring multiple NICs using DHCP, ensure that same DNS is used to resolve the IP addresses for all NICs.
- In case of multiple Resiliency Managers, you need to either use the same NTP server for configuration or ensure that the NTP servers are properly synchronized.
- Veritas Resiliency Platform supports Linux NTP server. You can also use a public NTP server If there is internet access to the appliance. It is recommended to use a pool (with odd numbers) of time resources for your NTP server. NTP server takes inputs from the available time sources and uses algorithms to find out the correct time. If there are even number of sources and they do not agree, then the algorithm of NTP server fails to make the right decision. Also, it is a better practice to use diversity of reference clocks. You can now configure NTP server using IPv4 or IPv6 address.
- Ensure that the NICs of the virtual appliances have a static MAC address. You can set a static MAC address for the appliance NICs using the virtual machine settings.
- While configuring the virtual appliances on the source data center, ensure that the IP and hostname of Resiliency Manager, IMS, and Replication Gateway can be resolved from the target data center and vice versa.

- In case of a Replication Gateway, make sure to attach an extra disk of at least 50 GB before configuring the Replication Gateway.

See [“Configuring the Resiliency Manager or IMS”](#) on page 439.

Configuring the Resiliency Manager or IMS

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

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

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

To configure the Resiliency Manager or IMS

1 Prerequisites:

See [“Prerequisites for configuring Resiliency Platform components”](#) on page 438.

- #### 2
- In any non-AWS environment, log in to the virtual appliance console or SSH using the following credentials:
 - **Username:** admin
 - **Password:** P@ssw0rd

Note: In Azure environment, the password would be the one provided during deployment.

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

See [“Password policies for Resiliency Platform virtual appliance”](#) on page 451.

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

- In AWS environment, do one of the following:
 - From Linux system:

Use SSH with the private key from the key-pair that you had selected while launching the instance in AWS. For example:

```
ssh -i private_key_file admin@ip_address_of_aws_instance
```

Ensure to modify the permissions for the private key file as 600 before using the file.
 - From Windows system:

Follow the documentation of AWS to connect to the Linux instance from a Windows system using PuTTY.
- 3** Accept the End User License agreement (EULA) to proceed with the configuration.
- 4** In the **Host Network Settings** section, you can configure the appliance network by using DHCP or static IP.

See [“Configuring network settings for Resiliency Manager”](#) on page 441.
See [“Configuring network settings for IMS”](#) on page 443.
- 5** In the **Appliance Settings** section, do the following:
 - Press the Enter key to confirm the use of an NTP server for configuring the date and time.
 - You are required to select the time zone. Follow the instructions as displayed on the virtual appliance console or SSH session to select the correct time zone.
 - Enter the FQDN or IP address of the NTP server. The appliance verifies the NTP server details. If there are any issues, details are displayed on the screen and you are prompted to enter the details again.

You can reset the timezone and NTP server at a later point of time using klish menu. Changing the system settings can affect the product functionality if incorrect values are set.
- 6** In the **Product Settings** section, the virtual appliance is configured as Resiliency Manager or IMS, depending upon the OVA file that you had selected for deployment.

- 7 After a successful product configuration, a message is displayed. If you have configured Resiliency Manager on the virtual appliance, a URL for the Resiliency Platform web console login is provided. You can type the URL in a web browser and log in to the web console.
- 8 If the bootstrap is in AWS environment, you must log in once again using the SSH key and set the admin user password. You need to use this password for subsequent logins to the console.

See [“About configuring the Resiliency Platform components”](#) on page 437.

Configuring network settings for Resiliency Manager

A Resiliency Manager needs to communicate with multiple entities within Veritas Resiliency Platform such as the IMS and another Resiliency Manager in the domain. To facilitate separate communication channels for these communications, Resiliency Platform 10.0 extends support for configuring Resiliency Manager with three Network Interface Cards (NIC).

The Resiliency Manager appliance is shipped with three NICs. You can configure these NICs to be used for the following three communications:

- For communication with Infrastructure Management Server (IMS)
- For communication with other Resiliency Managers and clouds too
- For communication with Product User Interface

If you do not plan to use three separate networks, you do not need to configure three NICs. You can configure one or multiple NICs based on your network layout. So, if you have only one network, you can configure only one NIC and associate all communications to go through the configured NIC.

Since IPv6 network support is provided from Resiliency Platform 3.3.1 version, you can configure the NICs using IPv4 and IPv6 addresses. For more information,

Note: The network configuration of the NICs performed during the bootstrap is final and you cannot edit the network configuration of the NICs at a later point of time.

To configure network settings for Resiliency Manager

- 1 In the **Host Network Settings** section, the bootstrap program first checks for network configuration of all the NICs of the appliance and prints the network details of all the NICs. You are prompted to continue the process of bootstrap. You can see a list of sections that need to be completed as a part of host network settings:
 - Host Network Settings for communication with Infrastructure Management Server

- Host Network Settings for communication to other Resiliency Managers
- Host Network Settings for communication with Product User Interface

2 For each section listed above, do the following:

- Enter the NIC to be used for the communication. The name and MAC address of NIC is displayed.
- The Bootstrap process checks if the NIC is already configured.
 - If the NIC is already configured then the NIC network configuration details are printed and you are prompted to confirm if you want to continue with the printed configuration.
 - If the NIC is not configured or if you do not want to use the existing NIC configuration, then you can choose to use either DHCP protocol or static protocol. In case of static protocol, you need to provide static network details such as IP address, prefix length.

Note: Ensure that appropriate subnet or virtual switch is assigned to the network adapter. Confirm this by matching the MAC address shown in the bootstrap with the one assigned to the virtual machine by the virtualization or cloud technology.

- Confirm if you want to add a static route. You need to set a static route for the interface only if you want the interface to reach a subnet that is different from all the subnets configured on this appliance and the default gateway is unable to communicate with that subnet. In this case, provide the subnet details (in CIDR format) at the input prompt. You can also set a static route to a host. In this case, provide the IP address of the host at the input prompt.
- Confirm if you are in Network Address Translation (NAT) environment and want to configure NAT when the NICs are configured in IPv4 networks only.

Note: Since NAT is not supported for IPv6 address, hence when you configure the virtual appliance using IPv6 address only and both IPv4 and IPv6 address, you are not asked about the NAT configuration. To know more about NAT support in Resiliency Platform

See [“About NAT support in Veritas Resiliency Platform”](#) on page 451.

You can add NAT gateway for communication between multiple Resiliency Managers, only if all the Resiliency Managers are not deployed in the same data center.

- 3 After successful configuration of the first NIC, confirm if you want to use the same NIC for the other communication channel. If you do not want to use the same NIC, perform step 2 for the other two communication channels.

Note: You can select one or more NICs for communication with Product User Interface. If you want to use multiple NICs to access product user interface, enter space separated values of the NICs.

- 4 Enter details for default router and then enter the DNS server details.
- 5 The details of **Host Network Settings** are displayed and you are prompted to confirm. Review the information. If any information is incorrect, choose **n** to go back to the networking inputs page and correct the details. Upon confirmation, the network is configured.

See [“Configuring the Resiliency Manager or IMS”](#) on page 439.

Configuring network settings for IMS

An IMS needs to communicate with multiple entities within Veritas Resiliency Platform such as the Resiliency Manager and the protected hosts. To facilitate separate communication channels for these communications, Resiliency Platform 10.0 extends support for configuring IMS with three Network Interface Cards (NIC).

The IMS appliance is shipped with three NICs. You can configure these NICs to be used for the following two communications:

- For communication with Resiliency Manager
- For communication with the gateways as well as with the hosts to be protected

Remaining one NIC is not configured during the bootstrap. Using the klish menu, you can configure that NIC later for communication between the appliance and any external entity. You can now configure this remaining NIC using IPv6 address with `nic-configuration set` command.

See [“Klish menu options for IMS”](#) on page 598.

If you do not plan to use two separate networks, you do not need to configure the two NICs. You can configure only one NIC or two NICs based on your network layout. So if you have only one network, you can configure only one NIC and associate the two communications to go through the configured NIC.

Since IPv6 network support is provided from Resiliency Platform 3.3.1 version, you can configure the NICs using IPv4 and IPv6 addresses. For more information,

Note: The network configuration of the NICs performed during the bootstrap is final and you cannot edit the network configuration of the NICs at a later point of time.

To configure network settings for IMS

- 1 In the **Host Network Settings** section, the bootstrap program first checks for network configuration of all the NICs of the appliance and prints the network details of all the NICs. You are prompted to continue the process of bootstrap. You can see a list of sections that need to be completed as a part of host network settings:
 - Host Network Settings for communication with Resiliency Manager
 - Host Network Settings for communication with Replication Gateway, Workload Virtual Machines, And Discovery Hosts
- 2 For each section listed above, do the following:
 - Enter the NIC to be used for the communication. The name and MAC address of NIC is displayed.
 - The Bootstrap process checks if the NIC is already configured.
 - If the NIC is already configured then the NIC network configuration details are printed and you are prompted to confirm if you want to continue with the printed configuration.
 - If the NIC is not configured or if you do not want to use the existing NIC configuration, then you can choose to use either DHCP protocol or static protocol. In case of static protocol, you need to provide static network details such as IP address, prefix length.

Note: Ensure that appropriate subnet or virtual switch is assigned to the network adapter. Confirm this by matching the MAC address shown in the bootstrap with the one assigned to the virtual machine by the virtualization technology or cloud technology.

- Confirm if you want to add a static route. You need to set a static route for the interface only. If you want the interface to reach a subnet that is different from all the subnets configured on this appliance and the default gateway is unable to communicate with that subnet. In this case, provide the subnet details (in CIDR format) at the input prompt. You can also set a static route to a host. In this case, provide the IP address of the host at the input prompt.
- Confirm if you are in Network Address Translation (NAT) environment and want to configure NAT when the NICs are configured in IPv4 networks only.

Note: Since NAT is not supported for IPv6 address, hence when you configure the virtual appliance using IPv6 address only and both IPv4 and IPv6 address, you are not asked about the NAT configuration. To know more about NAT support in Resiliency Platform

See [“About NAT support in Veritas Resiliency Platform”](#) on page 451.

You can add NAT gateway for communication between Resiliency Manager and IMS, only if IMS and the Resiliency Manager are not deployed in the same data center.

- 3 After successful configuration of the first NIC, confirm if you want to use the same NIC for the other communication channel. If you do not want to use the same NIC, perform step 2 for the other NIC.
- 4 Enter details for default router and then enter the DNS server details.
- 5 The details of **Host Network Settings** are displayed and you are prompted to confirm. Review the information. If any information is incorrect, choose **n** to go back to the networking inputs page and correct the details. Upon confirmation, the network is configured.

See [“Configuring the Resiliency Manager or IMS”](#) on page 439.

Configuring the Replication Gateways

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

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

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

To configure a Replication Gateway

- 1 Prerequisites:
 - See [“Prerequisites for configuring Resiliency Platform components”](#) on page 438.

- The Replication Gateway configuration requires an extra disk of at least 50 GB, to be used as a staging disk. You can attach this disk before configuration or during the configuration. The default disk size of 50GB lets you protect up to 8 virtual machines and each additional virtual machine requires a disk of 6GB size. You can increase the size of the disk by using `lvm` option of the klish commands:
- 2
- In any non-AWS environment, log in to the virtual appliance console or SSH using the following credentials:
 - **Username:** admin
 - **Password:** P@ssw0rd

Note: In Azure environment, the password would be the one provided during deployment.

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

See [“Password policies for Resiliency Platform virtual appliance”](#) on page 451.

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

- In AWS environment, do one of the following:
 - From a Linux client system:

Use SSH with the private key from the key-pair that you had selected while launching the instance in AWS. For example:

```
ssh -i private_key_file admin@ip_address_of_aws_instance
```

Ensure to modify the permissions for the private key file as 600 before using the file.
 - From a Windows client system:

Follow the documentation of AWS to connect to the Linux instance from a Windows system using PuTTY.
- 3
- Accept the End User License agreement (EULA) to proceed with the configuration.

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

See [“Configuring network settings for Replication Gateway”](#) on page 448.

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

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

- 6** In the **Product Settings** section, the virtual appliance is configured as Replication Gateway.
- 7** You are prompted to confirm if you want to enable FIPS for the Replication Gateway appliance.

See [“About FIPS enablement for Replication Gateway appliance”](#) on page 450.

Note: If you confirm to enable FIPS for the appliance, the appliance will be restarted after finishing the bootstrap process.

- 8** You are prompted to attach an extra disk to the appliance. If you have already attached the extra disk, press **Enter** to confirm. If you have attached more than one extra disk, all disks are listed and you need to select the extra disk that you want to use. If you have not already attached the extra disk, attach the extra disk and then confirm or select the extra disk to be used. Ensure that you attach a thick provisioned disk.

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

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

If the bootstrap is in AWS environment, you must log in once again using the SSH key and set the admin user password. You need to use this password for adding the gateway to the resiliency manager.
- 10 Add the Replication Gateway to an IMS using the Resiliency Manager console.

See [“Adding a Replication Gateway”](#) on page 111.

Configuring network settings for Replication Gateway

A Replication Gateway needs to communicate with multiple entities within Veritas Resiliency Platform such as the IMS, protected hosts, and peer Gateway. To facilitate separate communication channels for all these communications, Resiliency Platform 10.0 extends support for configuring Replication Gateway with multiple Network Interface Cards (NIC).

The Replication Gateway appliance is shipped with four NICs. Out of these four NICs, you can configure three NICs to be used for the following three communications:

- For communication with Infrastructure Management Server (IMS)
- For communication with peer Replication Gateway
- For communication with the virtual machines to be protected

Remaining one NIC is not configured during the bootstrap. Using the klish menu, you can configure that NIC later for communication between the appliance and any external entity. You can now configure this remaining NIC using IPv6 address with `nic-configuration set` command..

See [“Klish menu options for Replication Gateway”](#) on page 612.

If you do not plan to use three separate networks, you do not need to configure all the three NICs. You can configure only one NIC or two NICs based on your network layout. So if you have only one network, you can configure only one NIC and associate all three communications to go through the configured NIC. Likewise, if you have two networks, you can configure two NICs and associate appropriate communications to go through these two NICs.

Since IPv6 network support is provided from Resiliency Platform 3.3.1 version, you can configure the NICs using IPv4 and IPv6 addresses. For more information,

Note: The network configuration of the NICs performed during the bootstrap is final and you cannot edit the network configuration of the NICs at a later point of time.

To configure network settings for Replication Gateway

- 1 In the **Host Network Settings** section, the bootstrap program first checks for network configuration of all the NICs of the appliance and prints the network details of the NICs. You are prompted to continue the process of bootstrap. You can see a list of sections that need to be completed as a part of host network settings:
 - Host Network Settings for communication to Infrastructure Management Server
 - Host Network Settings for communication to Peer Gateways
 - Host Network Settings for communication to virtual machines to be protected
- 2 For each section listed above, do the following:
 - Enter the NIC to be used for the communication. The name and MAC address of NIC is displayed.
 - The Bootstrap process checks if the NIC is already configured.
 - If the NIC is already configured then the NIC network configuration details are printed and you are prompted to confirm if you want to continue with the printed configuration.
 - If the NIC is not configured or if you do not want to use the existing NIC configuration, then you can choose to use either DHCP protocol or static protocol. In case of static protocol, you need to provide static network details such as IP address, prefix length.

Note: Ensure that appropriate subnet or virtual switch is assigned to the network adapter. Confirm this by matching the MAC address shown in the bootstrap with the one assigned to the virtual machine by the virtualization or cloud technology.

- Confirm if you want to add a static route. You need to set a static route for the interface only. If you want the interface to reach a subnet that is different from all the subnets configured on this appliance and the default gateway is unable to communicate with that subnet. In this case, provide the subnet details (in CIDR format) at the input prompt. You can also set a static route to a host. In this case, provide the IP address of the host at the input prompt.
- Confirm if you are in Network Address Translation (NAT) environment and want to configure NAT when the NICs are configured in IPv4 networks only.

Note: Since NAT is not supported for IPv6 address, hence when you configure the virtual appliance using IPv6 address only and both IPv4 and IPv6 address, you are not asked about the NAT configuration. To know more about NAT support in Resiliency Platform

See “[About NAT support in Veritas Resiliency Platform](#)” on page 451.

You can add NAT gateway for communication between the peer gateways.

- 3 After successful configuration of the first NIC, confirm if you want to use the same NIC for other two communication channels. If you do not want to use the same NIC, perform step 2 for rest of the NICs.
- 4 Enter details for default router and then enter the DNS server details.
- 5 The details of **Host Network Settings** are displayed and you are prompted to confirm. Review the information. If any information is incorrect, choose **n** to go back to the networking inputs page and correct the details. Upon confirmation, the network is configured.

See “[Configuring the Replication Gateways](#)” on page 445.

About FIPS enablement for Replication Gateway appliance

Federal Information Processing Standards (FIPS) is a set of standards that describe document processing, encryption algorithms, and other information technology standards for use within the non-military government agencies and by government contractors and vendors who work with these agencies.

Resiliency Platform Data Mover lets you configure encryption of data over WAN (Wide Area Network). It uses the openssl library provided by the operating system vendor RedHat to encrypt the data before transmitting the data. Openssl library shipped with RHEL 6.6 and onwards is certified under the certificate numbers 2446 and 2447:

<http://csrc.nist.gov/groups/STM/cmvp/documents/140-1/1401val2015.htm>

Veritas Resiliency Platform provides 128 bit and 256 bit encryption schemes. When the operating system gets started with FIPS mode enabled, Veritas Resiliency Platform operates in the FIPS compliant mode.

You can enable the FIPS mode for a Replication Gateway during the bootstrap process:

See “[Configuring the Replication Gateways](#)” on page 445.

You can enable, disable, or view the status of the FIPS mode by using the `manage > fips` option of klish menu:

Note: You need to keep similar FIPS setting for the paired Replication Gateways.

Password policies for Resiliency Platform virtual appliance

To access Resiliency Platform virtual appliances, you need to set a password for the admin user. Following is the list of rules to set the password:

- Must be at least 8 characters long.
- Must contain at least one uppercase letter (A-Z), one lowercase letter (a-z), one numeric (0-9), and one special character such as @&%.

Note: Though special characters are allowed in a password, you cannot use a space, dollar sign (\$) or double quotes (") in a password. However, this is applicable only in case of Resiliency Platform Data Mover.

- Cannot contain the user name or its characters in reversed order.
- Cannot contain same character used consecutively for more than 2 times.
- Cannot contain 5 or more characters from the previous password.
- Cannot be the same as your previous 6 passwords.
- Can be changed after a minimum of 15 days since the last password change. The password can be changed only through klish menu.
- Expires in 90 days. you get an error message when you are not able to login using admin user credentials.
- 7 days before the password expiry date, a warning is provided to change the password. This warning is not displayed in the Resiliency Manager console. You get to see this warning only when you log in to the virtual appliance console or in the SSH session using the admin user credentials.
- Maximum 10 authentication attempts are allowed within a duration of 15 minutes. After this limit, the user ID and password gets blocked for next 60 minutes.
- In case you forget the admin password, you need to contact Veritas support. Veritas support can reset the admin password only if the virtualization technology or cloud environment supports the serial console.

About NAT support in Veritas Resiliency Platform

Network Address Translation (NAT) is a process in which one or more computers inside a private network are assigned a public address. NAT reduces the need for IPv4 public addresses and hides private network address ranges.

Resiliency Platform 3.3 provides support for NAT to enable communication from a private network to an external network. If there is a non-routable network between the source data center and the target data center, then You need to configure NAT only using IPv4 address for communication between Resiliency Platform appliances.

Considerations for configuring NAT

- Resiliency Platform supports NAT only for the communication between Resiliency Platform components deployed in different data centers. These components need to communicate with each other over public IP address. NAT can exist in both the data centers or in any one of the data centers.
- Resiliency Platform does not support NAT for communication between components deployed in the same data center. The components based in the same data center can communicate with each other over private IP address.
- If a communication channel has been setup in a way that private IP address of the Resiliency Platform components is accessible in the other data center, then you need not configure NAT during the bootstrap process. For example, if you have setup VPN for communication between on-premises data center to AWS datacenter, then you need not configure NAT for Resiliency Platform components .
- If there are multiple Resiliency Managers in one data center, then either NAT configuration should be done for all of them or for none of them. A mix of NAT configured Resiliency Managers and non-NAT configured Resiliency Managers within a single data center is not supported.

Following are the scenarios in which NAT configuration is required for Resiliency Platform virtual appliances:

Table 2-31 Scenarios in which NAT configuration is required

Communication between	Description	NAT configuration
Replication Gateway with Peer Replication Gateway	A Replication Gateway can communicate with the peer Gateway only over Public IP.	Required
IMS with Resiliency Manager	If an IMS reports to a Resiliency Manager within the same data center, then it can communicate over private IP.	Not required
	If an IMS reports to a Resiliency Manager deployed in another data center, then it can communicate only over public IP.	Required

Table 2-31 Scenarios in which NAT configuration is required (*continued*)

Communication between	Description	NAT configuration
Resiliency Manager with another Resiliency Manager	If there are multiple Resiliency Managers in one data center and no Resiliency Manager in another data center, then all of the Resiliency Managers within the data center can communicate with each other over private IP.	Not required
	If there are multiple Resiliency Managers in one data center and at least one Resiliency Manager in another data center, then all of the Resiliency Managers need to communicate with each other only over public IP. In this case, you need to enable NAT reflection (Loopback NAT) for the data center where multiple Resiliency Managers are located.	Required

Set up the resiliency domain

Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console. Refer to the below topics:

- See [“Getting started with a new configuration”](#) on page 453.
- See [“Adding an IMS ”](#) on page 456.
- See [“Adding a Replication Gateway”](#) on page 111.
- See [“Adding Google Cloud Platform data center”](#) on page 458.
- See [“Managing user authentication and permissions”](#) on page 463.
- See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Getting started with a new 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.

Table 2-32 Getting Started wizard

Wizard setup	Details
<p>1. Create or Join a Resiliency Domain</p>	<p>For a new Resiliency Platform deployment, select the option Create Resiliency Domain and supply a name for the domain. You can choose whether to allow collection of product usage information.</p> <p>Select Join Resiliency Domain option if you already have a resiliency domain set up and want to add another Resiliency Manager to the existing domain.</p> <p>Refer to See “Adding a Resiliency Manager to an existing resiliency domain” on page 455. Click Continue.</p>
<p>2. Set up the Resiliency Manager</p>	<p>Select Cloud Data Center? if this is a public cloud data center.</p> <p>Enter the geographical location of the data center in Data Center Location.</p> <p>Provide a user friendly name to the data center in Data Center Name.</p> <p>Provide a user friendly name to the Resiliency Manager in Resiliency Manager Name.</p> <p>Default entries are shown if the Resiliency Manager has external Internet access to determine the geographical location.</p> <p>Click Create.</p>
<p>3. Set up Authentication Domain</p>	<p>Optional.</p> <p>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.</p> <p>If you want to assign a different user as Super admin you must first set up an LDAP or Active Directory authentication domain.</p> <p>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.</p>
<p>4. Set up Users and Personas</p>	<p>Optional.</p> <p>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.</p> <p>The user with the Super Admin persona can add other users and groups and assign them personas later using the Settings page.</p>

Table 2-32 Getting Started wizard (*continued*)

Wizard setup	Details
5. Set up Cloud Configuration	<p>Optional.</p> <p>This step is enabled only if you select Cloud Data Center in the step 2.</p> <p>You can skip this step and add the cloud configuration later from the console.</p> <p>The wizard verifies the information you enter and notifies you if the information is invalid.</p>
6. Finish Getting Started	<p>You exit the Getting Started wizard. The Dashboard is displayed and from the Settings page you can complete any steps that you have skipped.</p>

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 as well as 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 and the Admin login credentials for an existing Resiliency Manager virtual appliance in the resiliency domain.
 - In case of multiple Resiliency Managers in a data center, ensure the following:
 - All the existing Resiliency Managers must be online in the data center where you plan to add a new Resiliency Manager.
 - For a cloud data center, it is recommended to have a minimum of three Resiliency Managers if you want to have multiple Resiliency Managers in the data center.
- 2 Log in to the web console on the new Resiliency Manager. The Getting Started wizard is displayed.

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

Enter the fully qualified host name of a Resiliency Manager in the domain you want to join, user name, and password for the Resiliency Manager, and click **Verify**.
- 4 In Set up Resiliency Manager, specify the data center location, the data center friendly name, and Resiliency Manager friendly name.**
- 5 Click Confirm & Continue.**
- 6 After the host name, user name, and password has been verified, the Resiliency Domain Name appears automatically. Select one of the following data centers and click Continue.**
 - **Create new data center**
 - **Select from existing data center**
- 7 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 456.
- 8 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.

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

- If the IMS is in AWS or Google Cloud Platform (GCP) cloud data center, ensure that the IMS has an IAM role with all the required permissions attached to it.
AWS: See “Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway” on page 389.
GCP: See “API permissions required for Google Cloud Platform” on page 460.
 - Information needed for adding the IMS:
Provide hostname.
The admin user credentials for the IMS virtual appliance. This information is optional and you need to enter only if the server is directly accessible. If the server is not directly accessible, you can still initiate the process of adding an IMS by entering only the data center, friendly name, and FQDN/IP address. In this case, you get a registration URL which you have to use after logging in to the virtual appliance console of the IMS that you want to add and then the IMS is added to the data center.
 - Ensure that the IP address and hostname of the IMS gets resolved from the Resiliency Manager.
- 2** Navigate to **Settings** (menu bar) > **Infrastructure** > **Details View** and then Select **+ Infrastructure Management Server**.
- You can also access this page from the **Quick Actions** menu > **Manage Asset Infrastructure**.
- 3** In **Add Infrastructure Management Server**, enter the information for the IMS and submit.
- Tips:
- You can select from a list of existing data centers or add a new data center.
- To specify a new data center, select **New** in the **Data Center** field, then specify the location and name. When entering the location, enter a form of location identifier, such as city, and the location list will populate with potential matches for you to select.
 - Enter a friendly name for the IMS.
 - Enter the FQDN or IP address of the server.
 - Enter the user name and the password. These two are optional information that you need to enter only if the IMS is directly accessible. If you provide this information, the IMS is immediately added to the data center. If you do not provide the username and password of the IMS, a registration URL is displayed on the screen. This URL is valid only for 30 minutes. If the URL expires, you need to regenerate the registration URL to complete the process.

Copy the URL string and then log in to the virtual appliance console of the IMS. In the klish menu, run the following command:

```
manage > configure ims_register
```

See “[Klish menu options for IMS](#)” on page 598.

You are prompted to provide the IMS registration URL. Enter the URL that you had obtained after initiating the process from Resiliency Manager console.

4 Verify that the IMS is successfully added.

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

Note: When the IMS is connected to a RM in it's own datacenter, that IMS is displayed alongside the RM in the same row on the UI. However, when the IMS is connected to a RM in another datacenter, that IMS is displayed in a row below the RM.

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.

Adding Google Cloud Platform data center

To access the cloud resources, you need to configure the cloud configuration. This information is used to validate the cloud configuration. When you add an Google Cloud Platform cloud data center, you enter a few details for the data center such as Project ID and Service Account like Client Emails, Private key if not autodetected. You need to enter details for region and bucket name also.

To add Google Cloud Platform data center

1 Prerequisite:

- Ensure that the Service Account which is used to add the Google Cloud Platform data center has the required API permissions.
- Ensure that the subnet of the NIC with which Resiliency Manager can communicate with GCP having private Google Access should be turned on or should have internet access.
- Ensure to enable Compute Engine, Cloud KMS, and Cloud Storage APIs.

See [“API permissions required for Google Cloud Platform”](#) on page 460.

- 2 Navigate to **Settings (menu bar) > Infrastructure > Details view**.
- 3 Select **Datacenter +**.
- 4 Enter the geographical location and the name of the data center.
- 5 Select **Is Cloud Datacenter** and then select **Google Cloud Platform** as cloud type. Click **Next**.
- 6 In the **Google Cloud Platform configurations** window, on **Service Account details for data center** panel select any one option to enter the following information:

- **Continue with associated account**
- **Provide another service account details**

If you select **Continue with associated account** option, provide following information:

- Enter the configuration name.
Provide a user friendly name to the configuration.
- **Project ID:** This is (auto-detected)
- **Service Account:**This is (auto-detected)
- **Region:** Select from the dropdown.
- **Bucket:** Select from the dropdown.
- You can select or unselect the **Discover the Shared VPC Resources** checkbox.

Note: If this checkbox is checked, at least one subnet from same region needs to be shared from host project with configured project. You need to configure few permissions for this checkbox. Refer [API permissions required for Google Cloud Platform](#)

If you select **Provide another service account details** option, provide following information:

- Enter the configuration name.
Provide a user friendly name to the configuration.
- **Project ID:** Enter the Project ID.
- **Client Email:** Enter the Client Email.
- **Region:** Select from the dropdown.

- **Bucket:** Select from the dropdown.
- You can select or unselect the **Discover the Shared VPC Resources** checkbox.

Note: If this checkbox is checked, at least one subnet from same region needs to be shared from host project with configured project. You need to configure few permissions for this checkbox. Refer [API permissions required for Google Cloud Platform](#)

- Enter Private Key. Click **Verify** to validate the private key. This key can be extracted from Service account json. Refer [Creating and managing service account keys](#)

7 Click **Submit** to complete the configuration.

[Managing cloud configurations](#)

[Refreshing cloud data center](#)

API permissions required for Google Cloud Platform

Following are the permissions required for the roles that you need to create for Resiliency Manager and IMS for recovery to Google Cloud Platform data center. If you are deploying your appliances through Marketplace, then refer to

[API permissions for deploying Resiliency Platform appliances using Google Cloud Platform through Marketplace.](#)

Table 2-33 API Permissions required for role for Resiliency Manager

Service name	Permissions
compute	compute.regions.list
storage	storage.buckets.list

Table 2-34 API Permissions required for role for IMS

Service name	Permission
cloudkms	cloudkms.cryptoKeyVersions.list
	cloudkms.cryptoKeys.list
	cloudkms.keyRings.list

Table 2-34 API Permissions required for role for IMS (*continued*)

Service name	Permission
compute	compute.addresses.list
	compute.diskTypes.list
	compute.disks.create
	compute.disks.createSnapshot
	compute.disks.delete
	compute.disks.list
	compute.disks.use
	compute.disks.get
	compute.firewalls.list
	compute.globalOperations.list
	compute.images.create
	compute.images.get
	compute.images.useReadOnly
	compute.instances.attachDisk
	compute.instances.create
	compute.instances.delete
	compute.instances.detachDisk
	compute.instances.get
	compute.instances.list
	compute.instances.setMetadata
	compute.instances.setTags
	compute.instances.start
	compute.instances.stop
	compute.machineTypes.list
compute.networks.list	

Table 2-34 API Permissions required for role for IMS (*continued*)

Service name	Permission
	compute.projects.get
	compute.regionOperations.list
	compute.snapshots.create
	compute.snapshots.delete
	compute.snapshots.list
	compute.snapshots.useReadOnly
	compute.subnetworks.list
	compute.subnetworks.use
	compute.zoneOperations.list
	compute.zones.list
	compute.addresses.useInternal
storage	storage.objects.create
	storage.objects.get

Table 2-35 Permissions required to discover Shared VPC for a role for IMS on the host project.

Service name	Permission
compute	compute.networks.list
	compute.firewalls.list
	compute.addresses.list

Note: To share subnets from the host project with configured project, you need to add the Service Account associated with IMS as **Principal** and provide a **Compute Network User role** on the shared subnets.

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 2-36 Process for setting up user authentication and permissions

Task	Details
Configure authentication domains	<p>You can add multiple authentication domains.</p> <p>See “About user authentication in the web console” on page 464.</p> <p>See “Configuring authentication domains” on page 469.</p> <p>See “Unconfiguring authentication domains” on page 474.</p>
Configure user groups and users	<p>Once you configure an authentication domain, you can configure user groups or users for Resiliency Platform from that authentication domain.</p> <p>See “Configuring user groups and users” on page 475.</p>
Assign permissions to groups and users	<p>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.</p> <p>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.</p> <p>See “About user permissions in the web console” on page 464.</p> <p>See “Predefined personas” on page 465.</p> <p>See “About limiting object scope for personas” on page 482.</p> <p>See “Assigning permissions to user groups and users” on page 476.</p> <p>You can also create custom personas.</p> <p>See “Adding custom personas” on page 478.</p> <p>See “Predefined jobs that can be used for custom personas” on page 479.</p>

Table 2-36 Process for setting up user authentication and permissions
(continued)

Task	Details
Configure Windows global user	To customize the static IP of Windows guest virtual machines in the VMware environment, you need to provide the administrator user name and password to log on to the Windows virtual machines. See “Configuring Windows global user” on page 482.

About user authentication in the web console

By default, the Admin user of the Veritas Resiliency Platform virtual appliance can log in to the web console with access to all views and operations.

The Admin user can configure authentication domains from external identity providers such as Active Directory (AD) and LDAP.

Once an authentication domain is configured, the Admin user can configure user groups and users for Resiliency Platform from that domain. These users can log in to the console with their domain login credentials.

All users and groups that are configured for Resiliency Platform have permission by default to view everything in the web console but not to perform any operations. Permissions for operations must be assigned separately by assigning the appropriate personas to users and groups.

It is recommended not to remove the default Resiliency Platform users or reduce the permissions of the default Resiliency Platform users.

If you change the password of a user who was configured to log on to the domain, you need to edit the configured domain and enter the new password for the user.

See [“Editing authentication domains”](#) on page 475.

See [“Managing user authentication and permissions”](#) on page 463.

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	<p>A role that has access to a predefined set of jobs (operations).</p> <p>The product comes with a set of predefined personas.</p> <p>See “Predefined personas” on page 465.</p> <p>You can also add custom personas.</p> <p>See “Adding custom personas” on page 478.</p> <p>See “Predefined jobs that can be used for custom personas” on page 479.</p> <p>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.</p>
Job	<p>A type of task (operation) that a user can perform.</p> <p>Examples:</p> <ul style="list-style-type: none"> Manage resiliency groups Manage assets Perform disaster recovery of resiliency groups
Object types and scope	<p>Each job can be performed on certain types of Resiliency Platform objects. Types of objects include data centers, resiliency groups, and virtual business services.</p> <p>When you assign a persona to a user or group, you define the scope of some jobs by selecting from available objects. For some jobs, the scope is the resiliency domain, which would be the entire scope of the product deployment.</p>

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 463.

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 2-37 Predefined personas and jobs

Persona	Description and scope	Jobs
Super admin	Can perform all operations on all objects in resiliency domain. Scope: Resiliency domain.	<ul style="list-style-type: none"> ■ All jobs ■ All objects in resiliency domain
Resiliency Platform admin	Manage Resiliency Managers and Infrastructure Management Servers (IMSS) and data centers. Manage assets. Manage user security settings and other product settings. Manage product updates. Scope: Resiliency domain.	<ul style="list-style-type: none"> ■ Manage enclosure assets ■ Manage server deployments ■ Manage product updates ■ Manage virtualization assets ■ Manage application host ■ Manage application cluster assets ■ Manage DR settings ■ Manage user security settings ■ Manage copy manager assets ■ Manage product settings ■ Manage access profiles ■ Manage cloud assets ■ Manage data mover assets
Resiliency Platform Deployment admin	Manage Resiliency Managers and Infrastructure Management Servers (IMSS). Can add an IMS to an existing data center. Manage product updates. Scope: Resiliency domain.	<ul style="list-style-type: none"> ■ Manage server deployments ■ Manage product updates

Table 2-37 Predefined personas and jobs (*continued*)

Persona	Description and scope	Jobs
Data Center admin	<p>Manage infrastructure pairing and manage assets of specified types.</p> <p>Scope: Specified data center.</p>	<ul style="list-style-type: none"> ■ Execute custom scripts ■ Manage cloud assets ■ Manage enclosure assets ■ Manage copy manager assets ■ Manage application cluster assets ■ Manage data mover assets ■ Manage DR settings ■ Manage application host ■ Manage virtualization assets ■ Manage access profiles
Resiliency Domain admin	<p>Create, update, and delete resiliency groups, virtual business services (VBSs), and resiliency plans and templates.</p> <p>Start/stop all resiliency groups and VBSs.</p> <p>Configure all resiliency groups for disaster recovery (DR).</p> <p>Perform rehearsal and DR operations: migrate, recover.</p> <p>Create, update, and delete resiliency plans and templates.</p> <p>Manage disaster recovery network settings.</p> <p>Scope: Resiliency domain.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Manage resiliency plans ■ Manage virtual business services ■ Manage resiliency plan templates ■ Manage resiliency groups ■ Recover resiliency group ■ Execute custom scripts ■ Rehearse resiliency group

Table 2-37 Predefined personas and jobs (*continued*)

Persona	Description and scope	Jobs
Resiliency Group admin	<p>Update and delete specified resiliency groups.</p> <p>Start/stop specified resiliency groups.</p> <p>Start/stop VBSs as long as the VBS contains only the specified resiliency groups.</p> <p>Scope: Specified resiliency groups.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Manage resiliency groups
Resiliency Group operator	<p>Start/stop specified resiliency groups.</p> <p>Start/stop VBSs as long as the VBS contains only the specified resiliency groups.</p> <p>Scope: Specified resiliency groups.</p>	Start/stop resiliency groups
VBS admin	<p>Create, update, and delete all virtual business services (VBSs).</p> <p>Start/stop all resiliency groups and VBSs.</p> <p>Scope: Resiliency domain.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Manage virtual business services
Resiliency Group Recovery admin	<p>Manage and perform disaster recovery of resiliency groups</p> <p>Start/stop specified resiliency groups.</p> <p>Start/stop or perform DR operations on VBSs as long as the VBS contains only the specified resiliency groups.</p> <p>Scope: Specified resiliency groups.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Manage resiliency groups ■ Recover resiliency groups ■ Rehearse resiliency groups

Table 2-37 Predefined personas and jobs (*continued*)

Persona	Description and scope	Jobs
Resiliency Group Recovery operator	<p>Start/stop specified resiliency groups.</p> <p>Perform disaster recovery on specified resiliency groups.</p> <p>Start/stop or perform DR operations on VBSs as long as the VBS contains only the specified resiliency groups.</p> <p>Scope: Specified resiliency groups.</p>	<ul style="list-style-type: none"> ■ Start/stop resiliency groups ■ Recover resiliency groups ■ Rehearse resiliency groups
Guest	<p>View all information in console.</p> <p>Assigned by default when user or group is configured for Resiliency Platform.</p> <p>Scope: Resiliency domain</p>	View all information
Resiliency Platform Assets admin	<p>Manage all assets such as enclosure, application, application cluster assets, virtualization, data mover, and cloud.</p> <p>Scope: Resiliency domain</p>	<ul style="list-style-type: none"> ■ Manage enclosure assets ■ Manage virtualization assets ■ Manage application host ■ Manage application cluster assets ■ Manage copy manager assets ■ Manage access profiles ■ Manage cloud assets ■ Manage data mover assets

See [“Managing user authentication and permissions”](#) on page 463.

Configuring authentication domains

By default, the Admin user on the Veritas Resiliency Platform virtual appliance can log in to the Resiliency Platform web console with access to all views and operations. The Admin user can configure authentication domains for Resiliency Platform from external identity providers so that other users can be authenticated for access to the console.

[To configure authentication domains](#)

[To edit authentication domains](#)

To configure authentication domains

1 Prerequisites

The fully qualified domain name (FQDN) or IP address and credentials for the LDAP/AD servers in the authentication domain. If you are configuring with IPv6 address, specify the hostname and not the IP address.

2 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Domains**

Note: You can also configure an authentication domain from the Getting Started wizard after setting up the Resiliency Manager and resiliency domain.

3 Click **+ Configure Domain**.

4 Select a data center and under **Specify server information for each data center**, enter the information for the server at that data center.

Repeat this step for other data centers in the authentication domain. When you select a different data center, the server information fields are cleared so that you can enter information for a different server, but the entries for the previous data center are remembered.

Note: If the same server is used for more than one data center, enter the same server information for each data center.

The remaining fields on the page apply to all data centers; fill these in as required.

See [“Options for authentication domain configuration”](#) on page 471.

Once you have entered information for all data centers, click **Next**.

5 Verify and complete the configuration:

In the **Domain Name** field, enter a friendly name for the authentication domain. If you configure the login screen to list domains, this name is listed.

Verify that the applicable data centers are listed. To make any changes, click **Back** to return to the previous screen. Once all is complete, click **Submit**.

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

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

To edit authentication domains**1** Navigate to the domain list as described in the procedure to configure authentication domains.**2** Select the authentication domain you want to edit and select the Edit option.

Note the following guidelines when editing:

- To add server information for a new data center, select the applicable data center and fill in the server information.
- To edit existing server information, select the applicable data center.
- To edit other information, you do not need to select each data center; the same information applies to all.
- If a data center no longer uses a separate server, replace the server information for that data center with the information for the server that is being used.
- To remove a data center from the authentication domain, use the Unconfigure option instead of the Edit option.

See [“Unconfiguring authentication domains”](#) on page 474.

See [“Managing user authentication and permissions”](#) on page 463.

Options for authentication domain configuration

The first page of the authentication domain configuration wizard is divided into two areas.

See [Table 1-48](#) on page 472.

See [Table 1-49](#) on page 472.

Server information by data center

You must specify the server information separately for each data center. When you select a different data center the server information fields clear so you can enter

the new information. If the same server is used for multiple data centers, enter the same information for both data centers.

Table 2-38 Server information by data center

Option	Description
Server (Mandatory)	Enter the fully-qualified host name or IP address of the LDAP server. If a secure session is configured with the LDAP server using SSL certificates, you must enter the fully-qualified host name that matches with the fully-qualified host name in the LDAP server certificate.
Port (Mandatory)	Displays the number of the port on which the LDAP server is configured to run. By default, this field displays the port number as 389. You can edit this port number, if required.
Connect using SSL/TLS	Select this check box to use the Secure Sockets Layer (SSL) or Transport Layer Security (TLS) certificates to establish a secure channel between the authentication broker and the LDAP server.
Certificate	Browse to the location of the trusted root CA certificate of the vendor that issued the LDAP server certificate.

Configuration options applicable to all data centers

The remaining fields apply to all data centers; fill these in as required.

Table 2-39 Configuration options applicable to all data centers

Option	Description
The authentication servers require me to log on.	Select this check box if the anonymous operations are disabled on the LDAP server and a bind user ID is required to proceed with configuring the LDAP-based authentication

Table 2-39 Configuration options applicable to all data centers (*continued*)

Option	Description
Bind User Name/DN	<p>Enter the complete Distinguished Name (DN) of the user that is used to bind to the LDAP server.</p> <p>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</p> <p>For example, you can provide the DN as Administrator@enterprise.domainname.com ENTERPRISE\Administrator</p> <p>For RFC 2307 compliant LDAP servers, specify complete bind DN.</p> <p>For example, cn=Manager,dc=vss,dc=veritas,dc=com</p> <p>The LDAP or the AD administrator can provide you the bind user name that you can use.</p>
Password	Enter the password that is assigned to the bind user name that you use.
Query Information:	
User (Mandatory)	<p>Under Query Information, enter the user name based on which the system detects the LDAP server-related settings. Ensure that the user name does not contain any special characters.</p> <p>The system determines the search base based on the user name that you specify in this field.</p>
Group	<p>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.</p> <p>The system determines the search base based on the group name along with the user name that you have specified.</p>

Once you have entered information for all data centers, click **Next**.

The **LDAP standard** panel is introduced to select and discover the LDAP schema for the configuration.

Table 2-40 Configuring attributes for LDAP standards

Option	LDAP standard description	Steps
RFC2307	Resiliency Manager uses LDAP schema RFC2307 standard to populate the required attributes and discover the information from LDAP server.	If you select this option and click Next , on Verify and Configure page the LDAP server uses RFC2307 schema for configuration.
Microsoft Active Directory	Microsoft Active Directory is an Active Directory server which uses LDAP protocol.	If you select this option and click Next , on Verify and Configure page the LDAP server uses Microsoft Active Directory schema for configuration.
Custom	The LDAP server uses the default schema as part of LDAP configuration.	<p>If you select this option, you need to provide following attributes for LDAP server and then click Next. The LDAP server will use the default schema present as part of the LDAP configuration.</p> <ul style="list-style-type: none"> ■ User Name: ■ User ID: ■ User description: ■ Group Name: ■ Group ID: ■ Group description: <p>Click Submit and Done.</p> <p>On Verify and Configure panel, the inputs provided are displayed.</p>

See [“Configuring authentication domains”](#) on page 469.

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 463.

Editing authentication domains

Using Resiliency Platform console, you can edit the configuration of an authentication domain. The newly introduced **Custom attributes panel**, allows to edit the the attributes for the LDAP schema.

To edit an authentication domain

1 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Domains**

2 Right-click the domain and select **Edit**.

3 Edit the values that you want to update and click **Next**.

4 Verify the domain configuration details and click **Submit**.

See [“Managing user authentication and permissions”](#) on page 463.

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 Prerequisite

The names of the user groups or users that you want to configure, as configured in the authentication domain.

2 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Users & Groups**

Note: To edit or remove an existing user or group, right-click the name in the list and select the appropriate option.

3 Click **Configure User or Group**.

4 Select the authentication domain.

5 Type the name of the user group or user. Click **Verify** so that the wizard can verify the name in the domain.

6 You can allow this user to access APIs by selecting the **Allow user to access Resiliency Platform APIs** option.

If you have access to APIs, you can generate an API access key and start using Resiliency Platform APIs. From version 3.5, the user which has persona with job **Manage API access key** can generate or revoke the API access key to other users as well.

Note: This option to provide API access is not applicable for user groups.

7 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 476.

See [“Managing user authentication and permissions”](#) on page 463.

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 464.

To assign permissions to user groups and users

1 Prerequisites

The users and groups must be added to Resiliency Platform before you can assign personas.

See [“Configuring user groups and users”](#) on page 475.

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 463.

Adding custom personas

Veritas Resiliency Platform provides a set of predefined personas with access to predefined jobs.

You can add custom personas by selecting from the predefined jobs.

For example, the predefined persona Resiliency Platform Admin includes the jobs for managing assets, managing security settings, and managing product settings. You could create an "Asset Manager" persona that includes only the managing assets job.

You cannot customize the Super admin persona, which has access to all jobs and all objects in the resiliency domain. You also cannot customize the Guest persona, which can view all information in the console.

To add custom personas

1 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Persona & Jobs > New Persona**

2 In the **New Persona** page, complete the following steps and submit:

- Assign a name and description to the custom persona.
- Select one or more jobs that you want to assign to the persona. The jobs are shown in categories depending on whether the scope is the entire resiliency domain or whether the scope can be customized to specific data centers or assets. Select the job from the appropriate category.

For example, if you want to assign a permission related to managing any resiliency group in the resiliency domain, select **Manage Resiliency Group** under the category of **For entire Resiliency Domain**. But if you want to limit permissions to specific resiliency groups, select **Manage Resiliency Group** under the category **For specific resiliency groups**.

See [“Predefined jobs that can be used for custom personas”](#) on page 479.

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 463.

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 2-41 Jobs for custom personas

Jobs	Description	Scope
View all information	View all information in console.	Resiliency domain
Manage user security settings	Manage authentication domains, users and user groups, personas.	Resiliency domain
Manage product settings	Manage general product settings such as alerts and notifications.	Resiliency domain
Manage server deployments	Edit Resiliency Manager information. Join a Resiliency Manager to a domain or leave a domain. Manage IMSs, including add, remove, edit, reconnect operations.	Resiliency domain
Manage product updates	Perform the operations available from the Product Updates page of the console.	Resiliency domain

Table 2-41 Jobs for custom personas (*continued*)

Jobs	Description	Scope
Manage service objectives	Activate service objectives from templates; manage activated service objectives.	Resiliency domain
Manage assets, by type: <ul style="list-style-type: none"> ■ Manage host assets ■ Manage virtualization assets ■ Manage data mover assets ■ Manage application cluster assets ■ Manage cloud assets ■ Manage copy manager assets ■ Manage enclosure assets ■ Manage access profiles 	Add, edit, or remove specific types of asset infrastructure	Resiliency domain or specific data centers
Manage resiliency groups	Create, update, and delete resiliency groups.	Resiliency domain or specific resiliency groups
Start/stop resiliency groups	Start and stop resiliency groups.	Resiliency domain or specific resiliency groups
Manage virtual business services	Create, update, and delete virtual business services (VBSs).	Resiliency domain or specific VBSs
Manage resiliency plans	Create, update, and delete resiliency plans. Note: The permission to execute a resiliency plan depends on a cumulative check on permissions for individual resiliency groups and VBSs in the plan. See "About limiting object scope for personas" on page 482.	Resiliency domain
Manage resiliency plan templates	Create, update, and delete resiliency plan templates.	Resiliency domain

Table 2-41 Jobs for custom personas (*continued*)

Jobs	Description	Scope
Execute custom scripts	Execute custom scripts as part of resiliency plans.	Resiliency domain or specific data centers
Rehearse resiliency groups	<p>Perform rehearsal and rehearsal cleanup.</p> <p>Note: There is no separate job to perform rehearsal of VBSs. If the assigned scope of this job includes all the resiliency groups in a VBS, Rehearsal operations can be performed on that VBS.</p> <p>See “About limiting object scope for personas” on page 482.</p>	Resiliency domain or specific resiliency groups
Recover resiliency groups	<p>Perform Recovery operations such as migrate, recover, resync.</p> <p>Note: There is no separate job to perform disaster recovery of VBSs. If the assigned scope of this job includes all the resiliency groups in a VBS, DR operations can be performed on that VBS.</p> <p>See “About limiting object scope for personas” on page 482.</p>	Resiliency domain or specific resiliency groups
Manage DR settings	Configure disaster recovery network settings, for example, mapping network settings for disaster recovery or replication gateway pairing.	Resiliency domain or specific data centers

See [“Predefined personas”](#) on page 465.

See [“Adding custom personas”](#) on page 478.

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 465.

See [“About resiliency groups with assets”](#) on page 528.

For example, you can assign one user the permission to perform operations on resiliency group RG1 and assign another user the permission to perform operations on RG2.

When planning persona assignments in which you select objects to limit the scope, take the following into account:

- Before you can select the objects such as resiliency groups to limit the scope of operations for a persona, the objects must first be created in Resiliency Platform.
- You need to plan for future maintenance on such limited scope personas. If more objects of that type are added later, you may need to edit existing personas for users or user groups in order to add permissions for the new objects.
- Keep in mind that operations on virtual business services (VBSs) that include multiple resiliency groups will fail unless the user performing the operation has permission for operations on all the resiliency groups in the VBS. The same limitation applies for workflow or resiliency plan operations that include multiple resiliency groups.
For example: a VBS is composed of RG1 and RG2. The operator has permission to perform operations on RG1 but not RG2. If they try to perform operations on the VBS, the operation will fail.

Configuring Windows global user

To perform IP customization on Windows virtual machine in VMware environment, Resiliency Platform requires any one of the following users:

- Domain administrator
- Local user who is part of administrator group on the Windows host where you want to perform IP customization
- Domain user who is part of administrator group on the Windows host where you want to perform IP customization
- UAC settings
- User Account Control: Run all administrators in Admin Approval Mode.

For more information on how to manage the settings, refer to Microsoft documentation.

For Windows Active Directory user, the Active Directory should be common for both, the primary and the recovery data center. The Active Directory should be configured before configuring the Windows Global user.

If a Windows virtual machine is part of a Windows Active Directory, ensure that you log on to the virtual machine at-least once using the Active Directory credentials. This is applicable only for VMware environment and if the recovery is on on-premises data center.

To configure Windows global user

1 Navigate



Settings (menu bar)

Under **Product Settings**, click **User Management > Windows Global User**

2 Click + **Configure User** to configure the user.

3 Select between Active Directory and Workgroup.

4 Enter the administrator user name and password. Click **Verify**.

For Workgroup user, enter user name as workgroupname\username. If the workgroup name is not customized then you can enter only the user name.

5 On successful verification, click **Next** and then **Finish** to submit the information.

See [“Network customization options”](#) on page 523.

Managing settings for alerts and notifications and miscellaneous product settings

See the following topics for information on configuring email and SNMP settings for notifications and reports, setting up rules for event notifications, configuring purge intervals, and changing telemetry settings.

See [“Adding, modifying, or deleting email settings”](#) on page 484.

See [“Adding, modifying, or deleting SNMP settings”](#) on page 486.

See [“Throttling the notifications”](#) on page 486.

See [“Downloading the MIB file”](#) on page 490.

See [“Setting up rules for event notifications”](#) on page 490.

See [“Adding, modifying, or removing Syslog server”](#) on page 491.

See [“Modifying the purge setting for logs and SNMP traps”](#) on page 493.

See [“Enabling or disabling telemetry collection ”](#) on page 493.

See [“Showing domains on login screen ”](#) on page 494.

See [“Downloading log files”](#) on page 494.

Adding, modifying, or deleting email settings

You can configure email settings to be used for different features, such as sending reports or receiving automatic email notifications of events. Veritas Resiliency Platform manages email notifications via Resiliency Managers. When Resiliency Managers are located in different geographical locations, the required email settings are likely different for each location. In that case, you add a separate email configuration for each location. You can send a test email to verify the settings. You can also modify or delete existing email configurations.

To add, modify, or delete email settings

1 Navigate



Settings (menu bar)

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

To add a new email configuration, select **Add Email Configuration**.

To modify or delete an existing one, right-click it and select **Modify** or **Delete**.

2 To add or modify an email configuration, go through the wizard pages and specify the options.

In **Server Information**, specify the following:

Name	Assign a unique name for the email configuration.
Email Server	Valid formats include: Fully Qualified Domain Name (FQDN), IP address, or, if the network handles DNS resolution for host names, a shortened host name. Examples: Host123, Host123.example.com, xxx.yyy.zzz.aaa.
SMTP Port	Enter the SMTP mail server port number. The default is 25.
From Email Address	Enter the email address to be shown as the sender of all the emails that are sent.
Friendly Email Name	Optionally, enter a name to be shown for the From address.
Send To	Enter the email address to which you want to send the email.

3 In **Security**, if you want to implement secure SMTP, select the checkbox and enter the user name and password.

4 In **Select Resiliency Managers**, select a Resiliency Manager in the data center location where these email settings apply.

5 In **Test Email Settings**, enter a valid email address, and enter a subject and message for the test email. Select **Send Test Email** to test your settings.

6 Review the information in the summary and submit

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Adding, modifying, or deleting SNMP settings

When an event takes place, you can configure SNMP traps to be sent. The traps are generated using SNMPv2 version. The community string is set to *public* for the generated traps. Resiliency Platform 10.0 enables support for IPv6 network. You can configure the SNMP traps using IPv6 address. If you want to receive the SNMP traps from Resiliency Platform, you can configure using the below mentioned steps:

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 (IPv4 or IPv6) or name of the host where the SNMP trap console is located. Example: Host123.example.com
SNMP Port	Enter the SNMP port number. The default port for the trap is 162.

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Throttling the notifications

Notifications in Resiliency Platform are raised when changes occur due to new operations that are performed or any configuration settings are changed. Some events in Resiliency Platform are short-lived and are auto-cleared when the event condition is successful. In Resiliency Platform, multiple notifications are raised due to which it is difficult to track the real issues. From version 4.0 of Resiliency Platform, throttling notifications feature is introduced that helps in suppressing the notification

(with specific duration) for some time. The suppressed notification is not seen in the notification list under **Logs > Notification** tab.

For example: If a vCenter server is disconnected, a notification is raised. You can throttle this notification for specific duration and is not seen under the notification list.

If a notification is throttled for some duration, and it gets cleared before the specified duration then this notification is not seen under the notification list under **Logs > Notifications** tab.

For example: When a resiliency group is created, "Replication state synchronizing" notification is raised. If this notification is throttled for 5 minutes, and meanwhile the resiliency group is created and is in ONLINE state before the specified duration, the notification gets cleared and is not seen under the notification list.

Following are the points to remember:

- If you have throttled any notification for a specific duration, it is not seen in the notification list until the specified duration.
- If the condition of the notification is still valid, it shall be displayed in the notification list.
- The associated throttled notifications may get removed if the source is deleted from the Resiliency Platform.
- Notifications can be throttled for infinite duration.

For some events where clearing event is not assigned to a particular source, such events has "indefinite time period" duration set by default.

For example: Risk Service, Reporting Service, Scheduling Service, rg.migrate.success, risk.notify.suppress.risk, vmware.vc.ims.refresh.success, etc.

Using Resiliency Platform console, you can add, edit, and remove the throttling notifications of the respective sources. You can group by object and notification topic from the list. The sources for which the notifications are generated in Resiliency Platform are:

Table 2-42 List of objects for which notifications are generated

List of objects
Discovery Host
Data center
ESX Server
Resiliency Group

Table 2-42 List of objects for which notifications are generated (*continued*)

List of objects
NetBackup Primary
Virtual Business Service
vCenter
Infrastructure Management Server
Cluster
Resiliency Manager
Replication Gateway

Table 2-43 List of services for which notifications are generated

List of services
Scheduling Service
Reporting Service
Risk Service

Disable throttling notifications

If you want to disable the throttling notifications, contact Veritas Support.

Notification Throttling Report

Notification Throttling Report displays all notifications that are currently throttled and are waiting to be raised. To view reports, **Reports menu > Inventory >**

Notification Throttling Report

See [“Viewing reports”](#) on page 580.

More Information:

See [“Add throttle notification”](#) on page 488.

See [“Edit throttle notification”](#) on page 489.

See [“Remove throttle notification”](#) on page 490.

Add throttle notification

To add a throttling notification in Resiliency Platform, perform following steps:

To add a new throttling notification

- 1 Navigate to **Settings (menu bar) > Product Settings > Alerts and Notifications**.
- 2 Click on the **Throttle Notifications** tab.
- 3 To add new throttle notification, click **+ Notifications**.
- 4 In the **Add Notification Topic Throttling Settings** panel, select the following:
 - Select the **Source** from the drop-down and click **Next**.
 - Select the notification source for the respective Source objects and click **Next**.
 - Select the check box for the notification topics for throttling of the Source from the list and click **Next**.
 - Set the duration for throttling the notification from the list and click **Next**. Using **Apply All** option, you can set the same duration to all the notifications at a time.
 - Click **Submit** and **Finish**.

The notification is added under the **Throttle Notifications** tab > Notifications list.

More Information:

See [“Throttling the notifications”](#) on page 486.

Edit throttle notification

To edit the throttling notification in Resiliency Platform, perform following steps:

To edit a new throttling notification

- 1 Navigate to **Settings (menu bar) > Product Settings > Alerts and Notifications**.
- 2 Click on the **Throttle Notifications** tab.
- 3 Select the specific **Source** from the list you wish to edit the notification.
- 4 In the throttling notifications list, select the notification you want to edit by clicking on the vertical ellipse and select **Edit** option. While editing the throttling notification, you can only change the duration of the notification.
- 5 Click **Next** and then **Finish** to save the changes.

See [“Throttling the notifications”](#) on page 486.

Remove throttle notification

To remove the throttling notification from the Resiliency Platform, perform the following steps:

To remove the throttling notification

- 1 Navigate to **Settings (menu bar) > Product Settings > Alerts and Notifications**.
- 2 Click on the **Throttle Notifications**.
- 3 Select the **Source** from the list you wish to remove the notification.
- 4 In the throttling notification list, select the notification you want to remove. Click on the vertical ellipses select **Remove** option. On the **Remove throttle notification** panel, you can **View Details** of the notification before removing it.
- 5 Click **Submit** to remove the throttle notification.

Note: If you are removing the throttling notification then the associated notification are also removed.

See [“Throttling the notifications”](#) on page 486.

Downloading the MIB file

You can download the management information base (MIB) file from the Resiliency Manager console. This MIB file defines the format of Veritas Resiliency Platform SNMP traps.

To download the MIB file

- 1 Navigate



Settings (menu bar)

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

- 2 Select **Download MIB file**.

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.

When a resiliency group is placed in maintenance mode, a notification is generated every 24 hrs. You can configure an SNMP or email alert for these notifications.

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 484.

See [“Adding, modifying, or deleting SNMP settings”](#) on page 486.

2 Navigate



Settings (menu bar)

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

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

To modify or delete an existing rule: Select the **Rules** tab, right-click the rule, and select **Modify** or **Delete**.

3 In **Configure Rule**, enter or modify the following:

Name	Enter a unique name for this rule.
Send emails to	Enter one or more email addresses separated by a comma
Send SNMP traps to	Optional
Select Notifications	Select one or more events that you want to be notified about

4 Select **Submit**.

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

Adding, modifying, or removing Syslog server

You can configure Veritas Resiliency Platform to share the Resiliency Manager logs with Syslog server using the Resiliency Platform console. You can configure Syslog server using the IPv6 address.

You can add multiple Syslog servers for a data center. But if you want to collect the audit logs on selected servers, then you need to add those to the Resiliency Platform first.

To add, modify, or remove Syslog server

1 Navigate



Settings (menu bar)

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

To add a new Syslog server, select **Add Syslog Server**.

To modify or delete an existing server, right-click the required server, and select **Edit** or **Remove**.

2 To add or modify Syslog configuration, specify the following:

Data Center Name	Select the data center whose logs you want to send to the Syslog server. Disabled for modify operation.
Syslog Server IP / name	Enter the Syslog server IP address (IPv4 or IPv6) or the name. Disabled for modify operation.
Port	Enter the port number.
Log Level	Select the log level from the following. <ul style="list-style-type: none">■ Critical: To share only the critical logs.■ Error: To share error and critical logs.■ Warning: To share warning, error, and critical logs.■ Informational: To share all the logs.
Send Audit Logs	Select if you want to share audit logs with the Syslog server.
Protocol	UDP is the default protocol.

3 Click **Next** and **Finish** to save the changes.

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Modifying the purge setting for logs and SNMP traps

By default, logs and SNMP traps are retained for two years. You can modify this purge setting.

To modify the purge setting for logs and SNMP traps

- 1 Navigate



Settings (menu bar)

Under **Product Settings**, click **Miscellaneous**

- 2 Under **Log Settings**, enter the new value for the purge setting, in months, and save the setting.

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

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.

The telemetry information is uploaded to Veritas telemetry collection servers if the Resiliency Manager has Internet connectivity.

To enable or disable telemetry collection

- 1 Navigate



Settings (menu bar)

Under **Product Settings**, select **Miscellaneous**

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

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Showing domains on login screen

You can set up the login screen to list the available authentication domains. By default, the domains list is not shown and the user must enter a fully qualified username, for example, `username@domain` or `domain\username`.

To show domains on login

1 Navigate



Settings (menu bar)

Under **Product Settings**, select **Miscellaneous**

2 Under **Login Settings**, select **Show domains list** and save the setting.

See [“Managing user authentication and permissions”](#) on page 463.

Downloading log files

Using the Resiliency Platform console, you can download the Resiliency Managers logs for troubleshooting. These support logs are collected by running the `support > loggather` command using the klish menu.

You can view the list of collected support log files for the selected Resiliency Manager. If you have multiple Resiliency Managers, you need to log on to each Resiliency Manager to view its logs. To download the log files, you must have *Resiliency Platform admin* persona with *Manage product settings* job assigned.

Click **Refresh** to view the latest logs.

To download the log files

1 Navigate



Settings (menu bar)

Under **Product Settings**, select **Miscellaneous**

2 Under **Support logs**, select the log file in the list and download.

See [“Managing settings for alerts and notifications and miscellaneous product settings”](#) on page 483.

Add asset infrastructure

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

- See [“Adding VMware virtualization servers”](#) on page 167.
- See [“Preparing host for replication”](#) on page 496.

Adding Hyper-V virtualization servers

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

To add Hyper-V virtualization servers

1 Prerequisites:

See [“Prerequisites for Microsoft Hyper-V virtualization discovery”](#) on page 495.

2 Navigate



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

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

3 Expand the data center > **Virtualization and Private Cloud** > **Hyper-V** tab.

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

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

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

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

Prerequisites for Microsoft Hyper-V virtualization discovery

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

Table 2-44 Requirements for Microsoft Hyper-V virtualization discovery

Type of discovery	Requirements
Virtual machine discovery	<ul style="list-style-type: none"> ■ The <code>VRTSsfmh</code> package must be installed on the Hyper-V Server (parent partition). This is done automatically by the IMS when you add the Hyper-V server to Resiliency Platform. ■ The Hyper-V role must be enabled. ■ The Windows Management Instrumentation (WMI) service must be running on the Hyper-V Server.
Exported storage discovery	<ul style="list-style-type: none"> ■ The Windows Management Instrumentation (WMI) service must be running on the guest.
Hyper-V	<ul style="list-style-type: none"> ■ Ensure that you add all the Hyper-V servers of the cluster to Veritas Resiliency Platform. ■ User provided during configuration of Hyper-V node of Failover Cluster should have full cluster permissions.

Preparing host for replication

To enable the replication in Resiliency Platform using Resiliency Platform Data Mover, you need to add the asset and prepare it for replication. Asset can be a physical machine or a virtual machine.

To prepare a host for replication

- 1 Ensure that you understand the use cases and prerequisites for adding hosts to an IMS.
 - See [“About adding host assets”](#) on page 503.
 - See [“Prerequisites for adding hosts”](#) on page 505.

Note: After you add a new data disk to a Windows host and attach it to an IDE controller, you need to initialize the disk. This needs to be done before performing the Prepare Host for Replication task.

- 2 Navigate to **Settings** (menu bar) > **Infrastructure** > **Details View**.
You can also access this page from the **Quick Actions** menu.
- 3 Go to the on-premises data center and click **Data Mover**.
- 4 Under **Resiliency Platform Data Mover**, click **Prepare host for replication**.
- 5 In the wizard, select the type of host to be added.

- Discovered Virtual Machines: Choose this option to select one or more virtual machines that are already discovered from your virtualization environment.
 - Non-discovered hosts (no-hypervisor configured or physical hosts): Choose this option to add virtual machines which are not yet discovered or which are physical hosts. You can import multiple hosts from a simple comma separated text file using this option.
- 6** Choose the desired configuration mode from the following options:
- Using VRP Console: You can configure a host using the username and password of the host. Resiliency Platform automatically installs the appropriate host agent packages on the host.
 - Using pre-deployed VRP Host Agent Packages: Choose this option if you do not have the username and password of the host or if you require full control for deploying the host agent packages using your tools and then configure the host without providing the username and password from the Resiliency Platform console. The VRP host agent packages can be downloaded from the IMS.
See [“Downloading Resiliency Platform host agent packages from IMS”](#) on page 501.
See [“Configuring Resiliency Platform host agent packages manually”](#) on page 501.

Note: Ensure that you choose the same IMS to add this host from which you chose to download the VRP host agent package bundle to install and configure on that host.

- 7** Select the IMS to which you want to add a host and click **Next**
- 8** Based on the selected Host Type and Configuration Mode specify the following inputs:

Host Type	Configuration mode	Required Inputs
Discovered Virtual Machines	Using VRP console	<ul style="list-style-type: none"> ■ Select one or more Virtual Machines and click Next. You can filter the virtual machines based on environment or on the virtualization server they are running on, or search them using their virtual machine name, family type, or host name. ■ Review the discovered Hostname for the virtual machine and modify it, if necessary. ■ Enter the username and password for each virtual machine in the table. See Prerequisites for adding host for information about using non- root user accounts to add Linux hosts. ■ Click  to delete a row. <p>See “Prerequisites for adding hosts ” on page 505. for information about using non-root user accounts to add Linux hosts.</p>
Discovered Virtual Machines	Using a pre-deployed VRP Host Agent	<ul style="list-style-type: none"> ■ Select one or more Virtual Machines and click Next. You can filter the virtual machines based on environment or on the virtualization server they are running on, or search them using their virtual machine name, family type, or host name. ■ Review the discovered Hostname for the virtual machine and modify it, if necessary.

Host Type	Configuration mode	Required Inputs
Non-discovered hosts Physical hosts	Using VRP console	<p>Type the host name and user credentials information in the table row. Use the following options to add information for multiple hosts.</p> <ul style="list-style-type: none"> ■ Use the <i>Add new row</i> option to add a blank table row. ■ Use the <i>Import from a text file</i> option to import information from a text file with comma separated values. To do this, select the desired text file and click Load host details. ■ Use the  icon to copy the details of the elected table row. You can edit the details of the newly added row. ■ Use the  icon to delete the desired row. <p>See “Prerequisites for adding hosts ” on page 505. for information about using non-root user accounts to add Linux hosts.</p>

Host Type	Configuration mode	Required Inputs
Non-discovered hosts or Physical hosts	Using a pre-deployed VRP Host Agent	<p>Type the host name in the table row. Use the following options to add information for multiple hosts.</p> <ul style="list-style-type: none"> ■ Use the <i>Add new row</i> option to add a blank table row. ■ Use the <i>Import from a text file</i> option to import information from a text file with comma separated values. To do this, select the desired text file and click Load host details. ■ Use the  icon to copy the details of the selected table row. You can edit the details of the newly added row. ■ Use the  icon to delete the desired row.

- 9 Once you are done with entering the data for all the hosts, click **Submit** to initiate the Prepare Host for Replication activity for each host.
- 10 Click the **Status** column value to navigate to the respective activity details view for a host to check if the activity is in progress or has failed.

The host is listed in the table for Replication Hosts on the Resiliency Platform Data Mover tab with Status column as Connected when the activity completes successfully.

When the recovery environment is AWS Cloud, the cloud storage and network drivers must be installed on the host in following conditions:

- Windows Paravirtual (PV) drivers to recover Windows hosts
- Xen block storage and networking drivers to recover SUSE Linux 11.4 hosts

If the compatible drivers are not installed already, the bundled versions of the required drivers are installed automatically when the hosts are added to resiliency group.

The bundled versions of the driver are located at:

- For SUSE Linux platform: `/var/opt/VRTSsfmh/spool/addons/store/
<VRTSitrptapversion_of_the_release>/AWSPVDriver`

Refer to the `Readme.txt` file in this directory for version information of the bundled drivers.

Managing multiple VMWare Virtual machines with same BIOS ID

If a VMWare virtual machine is already protected with Veritas Standalone Replication and has in-guest component then you cannot provision another virtual machine with the same BIOS ID as that of the protected virtual machine unless the protected virtual machine is upgraded to version 3.5.

Downloading Resiliency Platform host agent packages from IMS

You can download the Resiliency Platform host agent packages compressed bundle from IMS to which you intend to add the hosts. When you deploy and configure the downloaded Resiliency Platform host agent package on a host, you can prepare it for replication from the Resiliency Platform console without user name and password. To add the host in the Prepare Host for Replication wizard, you have to select the same IMS from which you downloaded the host agent package.

To download the Resiliency Platform host agent packages from an IMS

- 1 Navigate to **Settings** (menu bar) > **Infrastructure** > **Details View**.
You can also access this page from the **Quick Actions** menu.
- 2 Go to the on-premises data center and locate the desired IMS card.
- 3 Click **Vertical Ellipsis** menu on the IMS card and click **Resiliency Platform Host Agent Package**.
- 4 The popup displays the host agent download URL. Use the Copy button to copy the link to the clipboard.
Alternatively, you may manually select the link and copy it.
- 5 Browse to the URL and download the zip bundle.

Configuring Resiliency Platform host agent packages manually

You can manually install and configure Resiliency Platform host agent packages on a host before preparing it for replication using Resiliency Platform Data Mover. The host agent package zip bundle can be downloaded from an IMS.

The Resiliency Platform host agent bundle contains a configuration script which configures the host so that the specific IMS can communicate with the host without using the host credentials.

Before you begin configuring the Resiliency Platform host agent packages, download the Resiliency Platform host agent packages zip bundle from the IMS on the host and extract it at a temporary location, for example, `/tmp` on Linux host or `c:\temp` on Windows host.

See [“Downloading Resiliency Platform host agent packages from IMS”](#) on page 501.

To configure Resiliency Platform host agent packages manually on a host

- 1 Install the packages on required host.
 - On Linux host
 - Separate folders for all supported platforms are pre-created inside Linux folder. Go to folder depending on which platform the package is to be installed. Copy `VRTSitrptap-<version>.rpm` to a temporary folder (for example: `/tmp`) on Linux host.
 - Install `VRTSsfmh-<version>.rpm` on Linux host.
 - Install `VRTSitrptap-<version>.rpm` on Linux host.
 - On Windows host
 - Copy `VRTSitrptap-<version>.exe` from `x86_64` folder on Windows host.
 - Install `VRTSsfmh-<version>.msi` on Windows host.
 - Install `VRTSitrptap-<version>.exe` on Windows host.
 - Restart the Windows host so that the replication driver is loaded.
- 2 Run the bundled script to configure the host for communication from IMS.
 - On Linux host
 - Execute the `ConfigureHostFor-<IMS-Hostname/IP>.pl` script on Linux host as:

```
# /opt/VRTSsfmh/bin/perl /tmp/ConfigureHostFor-<IMS-Hostname/IP>.pl --configure
```
 - On Windows host
 - Execute the `ConfigureHostFor-<IMS-Hostname/IP>.pl` script on Windows host as:

```
C:\> "C:\Program Files\Veritas\VRTSsfmh\bin\perl.exe" C:\temp\ConfigureHostFor-<IMS-Hostname/IP>.pl --configure
```

Configuring password-less sudo privileges for user account on Linux host

You can add a Linux host to an IMS or move it to another IMS using a user account that does not belong to root group, provided that the user account has sudo privileges to run the following commands without requiring a sudo password:

- `/bin/rpm`
- `/opt/VRTSsfmh/bin/perl`

Refer to the sample sudo configuration entry below for a user account named `user1`:

```
user1 ALL=(root) NOPASSWD: /bin/rpm,/opt/VRTSsfmh/bin/perl
```

Note: Ensure that the sudo binary path is included in the system PATH environment.

If the user account is part of a group for which sudo configuration is defined, the corresponding sudo configuration section for that group should have an entry for that user account to run commands with NOPASSWD label.

Versions like Linux RHEL 6.8, RHEL 7.2 and Centos 6.7, by default enable the `requiretty` property in the sudo configuration. This property must be disabled to allow running commands using sudo without a password prompt. To disable the property, comment the line `Defaults requiretty` in the sudo configuration. A sudo user (in case of linux workloads) requires the sudo package version to be minimum 1.7.8.

About adding host assets

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

Table 2-45 Use cases for adding host assets

Use case	Details
Application discovery and management	<p>For discovery of supported applications on either physical systems or virtual machines, you must add the physical system or virtual machine as a host.</p> <p>Note: For the use case of discovering and managing virtual machines rather than applications, you must add the virtual machines as hosts.</p> <p>For discovery of a custom application, after you add the hosts, you must also add the application on the Assets page.</p>
VMware vCenter Server discovery (optional)	You can add a host to be used by the IMS for discovery of a VMware vCenter Server.
Hardware replication	<p>For storage array-based replication, you may need to install array-specific software on a host and add the host as a discovery host.</p> <p>More information is available on requirements for adding enclosures for array-based replication.</p>
Resiliency Platform Data Mover host	To manage and protect virtual machines using Resiliency Platform Data Mover, you need to add the virtual machines as hosts.
Manage physical machines	To manage and protect physical machines, you need to add the physical machines as hosts.

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

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

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

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

See [“Prerequisites for adding hosts”](#) on page 505.

Prerequisites for adding hosts

Before you add hosts to Veritas Resiliency Platform for discovery and monitoring by an Infrastructure Management Server (IMS), ensure that the following prerequisites are fulfilled.

General prerequisites for adding host assets:

- Ensure that the IMS can communicate with the host.
- Ensure that the time difference between the system clocks on the IMS and on the host is not more than 90 minutes. It is recommended to configure NTP on the virtual machine that needs to be secured. NTP should be configured in such a way that when the virtual machine is on the source data center, it remains in sync with the IMS on source data center. After migration, the virtual machine should be in sync with the IMS on the target data center.
- If a text file with comma separated values is used to add hosts, ensure that it uses the correct syntax.
- Ensure that the vCenter server and the hosts which are being prepared for in-guest protection need to be added to the same IMS and that they cannot be part of two different IMS.

Additional prerequisites for Linux systems

- To add a Linux host with the configuration mode option, *Configure using VRP Console*, use a user account that is a part of the root group or a user account that has password-less sudo privileges.
See [“Configuring password-less sudo privileges for user account on Linux host”](#) on page 503.
- Ensure that `openssh-clients` package is present in the system. Typically it is installed in the operating system by default.
- For Resiliency Platform Data Mover host, it is recommended to install and configure `ntpd`. It ensures that the system’s time remains in sync even after migration to the other site.
- For Resiliency Platform Data Mover host, ensure that `dmidecode` package is available in the system. Typically it is installed in the operating system by default. However, it may be not present in minimal OS installations.
- In order to install the host package while adding the Linux host, ensure that the `PasswordAuthentication` field is set to **yes** in the `/etc/ssh/sshd_config` file on the host.
- You need to remove all the stale network files from the below directory before adding the host.
 - **RHEL and Centos path:** `/etc/sysconfig/network-scripts`

- **SUSE:** `/etc/sysconfig/network`
- If a host is being added or prepared using a sudo user (in case of linux workloads), it requires the sudo package version to be minimum 1.7.8.
- For Suse 15, ensure that `insserv` package is installed in the system. Typically, it is installed in the operating system by default.
- In order to persist the iptable rules, ensure that the iptables service is running and the `IPTABLES_SAVE_ON_STOP`, `IPTABLES_SAVE_ON_RESTART` and `IPTABLES_SAVE_COUNTER` directives are set to **yes** in the iptables config file.

Additional prerequisites for Windows systems

- To add a Windows host, it is recommended to use a domain user account with local administrator privileges.
If you cannot use a domain user account with local administrator privileges, you have an option to use an Administrator user or a user in local administrator group with required prerequisites.
See [“User account required for adding a Windows host”](#) on page 507.
- Ensure that you have appropriate User Access Control (UAC) settings for the user that is used for adding the Windows host.
See [“UAC settings required for adding a Windows host”](#) on page 507.
- The Windows Management Instrumentation (WMI) service must be running.
- Ensure that the password for the host does not contain a double quotes character.
- After you add a new data disk to a Windows host and attach it to IDE controller, you need to initialize the disk. This needs to be done before adding the host.
- If you have McAfee antivirus already installed on the virtual machines, you need to disable the other encryption options before adding the host.

Additional prerequisites for Oracle Discovery Hosts

- VRP supports Redhat Enterprise Linux hosts as Oracle Discovery Hosts.
- You have to install Oracle Client (Net) on the host. Oracle Net on the host must be configured by “oracle” user such that the `tnsnames.ora` points to all the Oracle instances to be discovered and managed. You can use the Oracle Net Configuration Assistance (`netca`) to configure the `tnsnames.ora` file.
- If you are configuring Oracle RAC, ensure that all RAC nodes are configured in the `tnsnames.ora` file. Make sure that `sqlplus` command on the discovery host is able to reach all the Oracle database instances individually.

- If the Oracle Database is replicated using Oracle DataGuard, make sure that you add only the database that is local to the Datacenter in the tnsnames.ora. You should use a different Discovery Host in the remote Datacenter for managing the remote Database.
- You may use the same discovery host to manage the multiple Oracle databases in the same datacenter. The Oracle Net has to be configured to be able to communicate with all the instances of all these databases.

See “[About adding host assets](#)” on page 503.

User account required for adding a Windows host

To add a Windows host, you need to have any one of the following credential sets:

Table 2-46 Credentials required for adding a Windows host

Windows host being added	Prerequisite
Domain user	Must have local administrator privileges on the Windows host being added.
Administrator	None
User in local administrators group	Add a registry entry to disable the remote restriction on the Windows host being added: Microsoft documentation Once the Windows host is added, you can enable the remote restriction again on that host.

UAC settings required for adding a Windows host

Following is the list of User Account Control (UAC) settings required for adding a Windows host in Veritas Resiliency Platform.

Table 2-47 UAC settings required for adding a Windows host

UAC Policy	Local administrator user	Domain user with administrator privileges	Domain user with local administrator privileges
Admin Approval Mode for Built-in Administrator Account	Disabled	Enabled	Enabled

Table 2-47 UAC settings required for adding a Windows host (*continued*)

UAC Policy	Local administrator user	Domain user with administrator privileges	Domain user with local administrator privileges
Allow UIAccess applications to prompt for elevation without using the secure desktop	Enabled	Disabled	Disabled
Behavior of the elevation prompt for administrator in Admin Approval Mode	Prompt for credentials on secure desktop	Prompt for credentials on secure desktop	Prompt for credentials on secure desktop
Behavior of the elevation prompt for administrator for standard users	Prompt for credentials	Prompt for credentials	Prompt for credentials
Detect application installations and prompt for elevation	Enabled	Enabled	Enabled
Only elevate executables that are signed and validated	Enabled	Enabled	Enabled
Only elevate UI Access application that are installed in secure locations	Enabled	Enabled	Enabled
Run all administrators in Admin Approval Mode	Enabled	Enabled	Enabled
Switch to the secure desktop when prompting for elevation	Disabled	Disabled	Disabled
Virtualize file and registry write failures to per-user locations	Enabled	Enabled	Enabled

Infrastructure Pairing

For recovering assets to the respective cloud platform, you have to do following infrastructure pairing. Refer the following topics:

-
- See [“About network objects”](#) on page 509.

- See [“Network pairs for recovering virtual machines to Google Cloud Platform \(GCP\)”](#) on page 512.

About network objects

Resiliency Platform discovers and displays information about layer 2 and layer 3 networks for the discovered assets.

Layer 2: The second level in the seven-layer OSI reference model, is used to transfer data between adjacent networks in a WAN or LAN environment. This layer is also known as Data Link Layer.

Examples of layer 2 networks: Port group/VLAN, vSwitch, cloud network and cloud subnet if the target data center is in cloud.

Layer 3: The network layer in the OSI reference model. mainly include routing and forwarding, as well as internetworking, addressing, packet sequencing, congestion control and further error handling.

Examples of layer 3 networks: Subnets and cloud subnets.

Resiliency Platform discovers and displays information about Layer 2 and Layer 3 networks in your datacenter. For VMware datacenter, Resiliency Platform 4.0 also supports NSX-T type networks present in vCenter server. It discovers and displays information about NSX-T provisioned L2 network objects from vCenter like Opaque switch, Opaque network and virtual distributed switch 7.0. NSX-T transport zones were discovered as vSwitches and NSX-T segments are discovered as port groups/VLAN.

Note: While adding transport zone in NSX-T Manager, either select NVDS or VDS for all the hosts.

For cloud technologies like AWS, Azure, vCloud Director and Google Cloud Platform, cloud subnets serve the purpose for both the layer 2 and layer 3 networks in the network pair. Network objects like private cloud subnet and private cloud network are listed under **Network Types** drop down irrespective of the cloud data center configuration.

You can manually add subnets, VMware port group/VLANs and Hyper-V VLANs that are not discovered in a data center. You cannot add vSwitches and cloud networks, if they are not discovered. Adding subnets using IPv6 address and pairing them across data centers is supported. Either while adding the network objects or while editing the discovered network objects, you need to choose a purpose. Purpose can be Production or Rehearsal.

See [“About Purpose”](#) on page 511.

For mapping the purpose of the network objects as Production or Rehearsal,

Network objects like private cloud subnet and private cloud network are listed under **Network Types** drop down irrespective of the cloud data center configuration.

A network pair is created using network objects across data centers. The network pairs should be defined before a resiliency group is created. Depending on whether the participating networks in the pair are layer 2 or layer 3 networks, the pair can be used for connecting the assets to the networks, or for assisting the customization of the IP addresses in the target network. The create network pair operation eliminates the need to manually connect each asset to a network at the target data center. For example, port group/VLAN to cloud network and subnet to port group/VLAN.

When the resiliency group is created, the network objects in the network pairs are evaluated by the Resiliency Platform. The CIDR (Classless Inter-Domain Routing) information from the network object is used to automatically calculate the IP address for the applicable assets in the respective target network if any of the below mapping is done:

- Subnet to subnet
- Subnet to cloud subnet
- Cloud subnet to cloud subnet
- Private Cloud Subnet to Private Cloud Subnet

The layer 2 network object pairs are mandatory to be defined for recovery to cloud environments, private cloud environments and recovery of physical machines to VMware environment. This mapping is optional for recovery from on-premises to on-premises environment. If the mappings are not defined for recovery to on-premise environment, then the virtual machine NICs are not connected to any network.

A layer 3 network pair is optional. If it is defined, the IP address for the asset is calculated based on the target subnet CIDR and can be further customized. If the network pair is not defined, then the IP address for the adapter gets assigned in one of the following ways:

1. If the IP customization option is checked, user must enter the IP address (IPv4 or IPv6 address depending upon the network configured) that needs to be assigned to the virtual machine NIC.
2. If the IP customization option is not checked, a DHCP (Dynamic Host Configuration Protocol) IP address gets assigned to the virtual machine adapter if the target technology supports it. (For example: cloud environments).
3. If the IP customization option is not checked, for on-premises to on-premises recovery, the virtual machine adapter IP settings are not changed.

When you perform a migrate, recover or rehearsal operation on a resiliency group, the Resiliency Platform evaluates the network pairs that have the layer 2 network objects and gets connected to the expected target network.

Using Resiliency Platform console, you can create network groups of cloud subnets for AWS cloud data center and port group/VLAN for VMware environment only.

In case of Google Cloud Platform, if shared subnets are discovered, the host project name is appended to the VPC name to distinguish the shared subnets from another project.

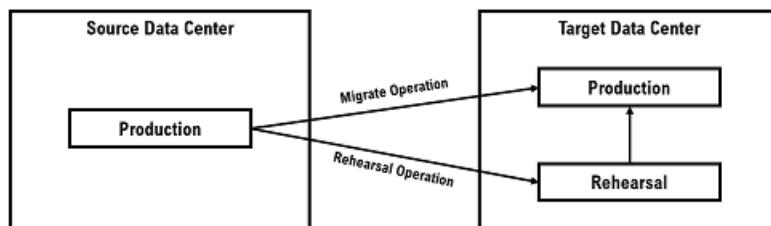
When you clone your virtual machines, ensure that you assign the appropriate host name and IP address to the cloned virtual machines.

About Purpose

When the corresponding network objects are involved in recovery tasks such as migrate, recover, or resync, then set the purpose as Production on those network objects. Similarly, when the corresponding network objects are involved only in Rehearsal operations, then set the purpose as Rehearsal on those network objects.

To perform recovery operations such as migrate, recover, or resync, you need to pair the network objects with a purpose as Production from source data center with network objects with a purpose as Production in the target data center. IPv4 network objects (for example, subnets) can be paired with only IPv4 network objects, and IPv6 network objects can be paired with only IPv6 network objects.

When you want to perform the rehearsal operations, you can map the network objects with purpose as Production with network objects (corresponding to rehearsal networks) with purpose as Rehearsal in the target data center. For example, you can map the AWS Cloud Subnet (Production as purpose) to AWS Cloud Subnet (Rehearsal as purpose). Following figure explains the mapping across the data center:



The create network pair operation eliminates the need to manually connect each virtual machine to a network at the recovery data center.

After you have paired the networks successfully, the target networks and the IP addresses are computed programmatically, and applied to the virtual machines.

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

For example:

In your source data center you have two subnets with Production and rehearsal as the purpose. The Production subnet is mapped with the Rehearsal subnet. A similar setup is present on your target data center. Both the subnets having the purpose as Production are paired.

Now when you perform the migrate operation, since the subnets having purpose as Production are paired, the virtual machine is migrated with appropriate network settings.

For example, a virtual machine has IP address 10.20.30.40 and is part of subnet 10.20.30.0/24 on source data center. This subnet is paired with another subnet 10.20.50.0/24 in the target data center. Hence when the virtual machine is migrated, its IP address is automatically changed to 10.20.50.40 on the target data center.

In the target data center, we have mapped the Production subnet to the rehearsal subnet. Hence when you perform the rehearsal operation, the rehearsal virtual machine is mapped on the rehearsal subnet.

For example, in your source data center you have two subnets (ProdSN1 10.20.30.0/24) with Production as the purpose and (ProdSN2 10.20.40.0/24) with rehearsal as the purpose. The Production subnet is mapped with the Rehearsal subnet. A similar setup is present on your target data center (RecovSN1 10.20.50.0/24 with Production as the purpose and (RecovSN2 10.20.60.0/24). Both the subnets having the purpose as Rehearsal are paired.

Note: Rehearsal subnet should be configured such that it is isolated from the production network.

You may want to perform Rehearsal operations in only one of the data centers. In that case, you can have network objects (for example, subnet) with purpose as Rehearsal in only that data center. If you map network objects with purpose as Rehearsal later, you need to edit the resiliency group with **Edit Configuration** or with **Customize Network** option.

Network pairs for recovering virtual machines to Google Cloud Platform (GCP)

You need to consider the below mentioned points before creating network pairs for recovering or migrating virtual machines to GCP :

- You can only configure a network interface when you create an instance.

- You cannot delete a network interface without deleting the instance.
- Each network interface configured in a single instance must be attached to a different VPC network, and each interface must belong to a subnet whose IP range does not overlap with the subnets of any other interfaces. The attached network can be a standalone VPC network or a Shared VPC network.
- For a workload multiple subnets having single VPC cannot be provisioned on GCP.
- For a workload multiple NICs in same subnet cannot be provisioned on GCP
- For shared subnets discovery, the host project name is appended to the VPC name to distinguish the shared subnets from the another project.

Figure 2-7 Overview of network structure for recovery to Google Cloud Platform

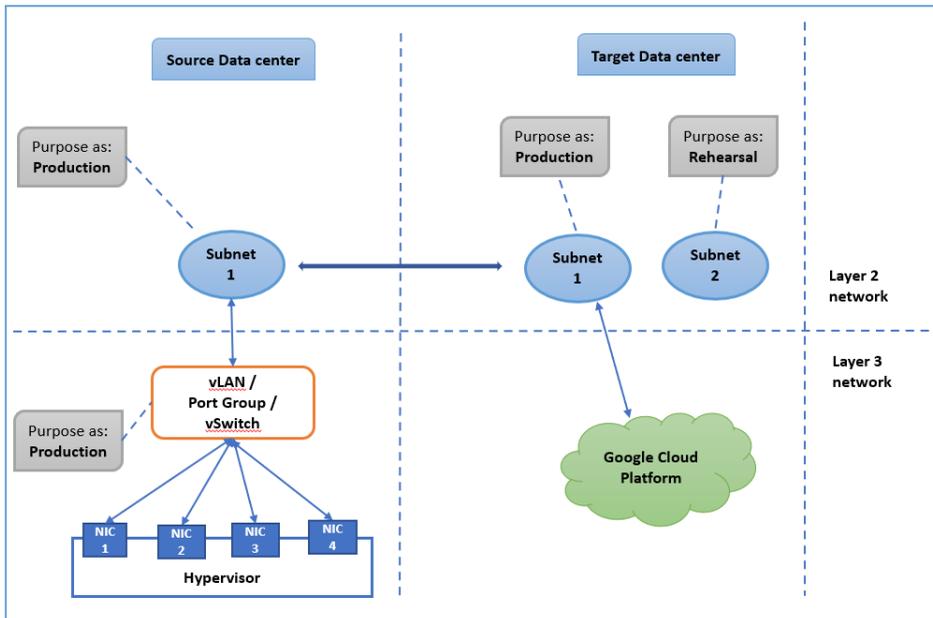


Table 2-48 Network Mapping types for Google Cloud Platform

Virtualization technology	Cloud environment	Network Mapping Type	Mandatory	Purpose
VMware	Google Cloud Platform	Subnet to Cloud Subnet	Yes	To connect the VNIC of the recovered virtual machine to the cloud network.
Hyper-V	Google Cloud Platform	Subnet to Cloud Subnet	Yes	To connect the VNIC of the recovered virtual machine to the cloud network.

[Creating network pairs between source and target data centers](#)

Creating network pairs between source and target data centers

The create network pair operation eliminates the need to manually connect each virtual machine to a network at the target data center. After you have paired the networks successfully, the target networks and the IP addresses are computed programmatically, and applied to the virtual machines.

Adding subnets using IPv6 address and pairing them across data centers is supported. You can pair subnets of same networks, that means pairing of subnets can be done for:

- IPv4 subnet to IPv4 subnet mapping
- IPv6 subnet to IPv6 subnet mapping

IPv6 network support is not applicable for cloud environments.

This step is a prerequisite for cloud recovery data center, if you want to override the default IP settings and customize the IP addresses when the virtual machines starts.

The mapping requirement depends on the target data center and is required to be done before you perform any disaster recovery operation. Before you create and map the network objects, ensure that all the assets are configured in the source as well as in target data center.

To create network pairs between production and recovery data centers

1 Navigate



Infrastructure Pairing (navigation pane)

2 Do one of the following:

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

3 In the **Network Mapping** page, select the data centers from source and the target data centers drop down to view the networks which are configured according to the vCenter or Hypervisor added .

Note that: Once you have configured hypervisor or cloud configuration then you will see the networks will be seen if they are discovered.

Networks Types options from the drop down will be displayed based on the virtualization technology or the cloud data centers configured.

4 If you want to view all the network objects, select the checkbox of **Show All Networks**. If you want to view only the network objects with purpose as **Production** , uncheck the checkbox.

5 Select the network objects from the lists and click **Move selected**.

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

Create resiliency groups

After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery. Refer following topics:

- Configuring a resiliency group for basic monitoring
- Prerequisites for configuring resiliency groups for recovery to Google Cloud Platform
- Configure resiliency groups for recovery to Google Cloud Platform

Configuring a resiliency group for basic monitoring

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

There are two types of pre-activated service objectives:

- Monitor assets - provides only monitoring, start, and stop operations
- Recover hosts - provides recovery operations as well as the start and stop operations

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

To manage assets for basic monitoring

1 Prerequisites

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

2 Navigate



Assets (navigation pane) **Unmanaged** tab > **Manage & Monitor Assets**

You can also launch the wizard from the **Overview** tab.

3 Select the assets:

- Select **Host** as the asset type, select the data center, type, and other filters as needed to display a list of assets.
- Drag and drop the selected assets to **Selected Instances**.

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

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

6 Supply a name for the resiliency group.

7 Verify that the new resiliency group is added to the **Resiliency Group(s)** tab.

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

Note: If the instances are created from BYOS image or there is licensing issues with instances, then Resiliency Platform operations may fail.

See [“About resiliency groups with assets”](#) on page 528.

Managing virtual machines for remote recovery (DR) to Google Cloud Platform

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

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

To manage virtual machines for remote recovery in Google Cloud Platform (GCP)

1 Prerequisites

See [“Prerequisites for configuring VMware virtual machines for recovery to Google Cloud Platform”](#) on page 206.

See [“Prerequisites for configuring Hyper-V virtual machines for recovery to Google Cloud Platform”](#) on page 518.

2 Navigate

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

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

3 Select the assets:

- Select the **Host** as Asset Type and select the other filters as needed to display a list of virtual machines.
- Drag and drop virtual machines or select the asset and click **Next**.

4 Next page displays the environment for the selected assets on **Review Environment** panel.

For resiliency group consisting of VMware virtual machines, ensure that each resiliency group is mapped to only one ESX cluster.

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

6 Select the target (recovery) data center on **Select Target Datacenter** page.

7 Review the information on the **Configure Resiliency Platform Data Mover** page and click **Next**.

- 8 Select the Replication Gateway pairs on **Select Replication Gateway pair** page.
- 9 Select the volume type.
See [“Volume type selection options”](#) on page 519.
- 10 In the **Confirm Resiliency Platform Data Mover Details** panel, verify the Replication Gateway pair selection and the asset information.
- 11 On the **Customize panel** perform the actions.
See [“Customize panel for Google Cloud Platform”](#) on page 520.
- 12 On the **Network Summary** panel, verify the asset name, MAC address and IP address of the assets in the resiliency groups and click **Next**.
- 13 Complete the network customization steps for the virtualization technology on the **Customize Network** panel.
See [“Network customization options”](#) on page 523.
- 14 Use the **Customize System Generated workflows** panel to enable manual intervention at predefined points during the Migrate and Recover operations on both the data centers. This is an optional step. See [“About manual intervention”](#) on page 525.
- 15 Verify the summarized information and enter a name for the resiliency group and click **Submit**.

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

Prerequisites for configuring Hyper-V virtual machines for recovery to Google Cloud Platform

- Before you run the wizard to configure disaster recovery protection for a resiliency group of Hyper-V virtual machines, ensure that you have met the following prerequisites for virtual machine configuration:
- If you add new disks, ensure that they are visible from the guest operating system.
- Ensure that the integrations services are installed and running inside the virtual machines.
- Ensure that you have disabled the 'Quick removal' policy for disks in Windows Server 2008 R2 and disabled the 'write-cache' policy for disks in Windows Server 2012 R2.

- When protecting Hyper-V virtual machines that use Dynamic Memory, ensure that the Maximum RAM value is set to a reasonable value. This value is used to assign the memory size to the virtual machines on the target (recovery) site. If Maximum RAM is set to a very high value then while configuring the resiliency group for remote recovery, the operation may fail.
- Ensure that the Replication Gateways have sufficient storage to handle the replication for the planned number of protected virtual machines. Both the on-premises gateway and the cloud gateway must have external storage equivalent to 6GB for each asset protected by the gateway pair. For example, if a gateway pair supports 10 virtual machines, the on-premises gateway and the cloud gateway must each have 60 GB of external storage.
- If the status of the virtual machine on the recovery data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.
- A Windows host is already added to resiliency domain using 'Prepare Host for Replication' operation. If you now add a non-initialized data disk to the host on IDE controller, you need to initialize the disk, reboot the host and refresh Hyper-V server before you can protect the host for DR.

See [“Managing virtual machines for remote recovery \(DR\) to Google Cloud Platform”](#) on page 517.

Volume type selection options

This panel is displayed when you are configuring your virtual machines for recovery to AWS or Google Cloud Platform.

Table 2-49 Volume type options available for AWS

Options	Description
Volume Type	Select the volume type for the disks. You can apply the selections either to all the virtual machines or customize for each.
IOPS (This field is enabled if Provisioned IOPS SSD option is selected)	Enter the IOPS required if the volume type is Provisioned IOPS SSD. Note: Refer to AWS documentation for more information on IOPS permitted for specific volume type and size.
KMS Encryption Key	Select the KMS Encryption Key, if you want to create encrypted volume in AWS, before selecting KMS key. Ensure that KMS key has all the required permissions. Note: Refer to AWS documentation for KMS key permissions and IAM role attached to IMS.

Table 2-49 Volume type options available for AWS (*continued*)

Options	Description
Availability Zone	The availability zones which are listed are based on the Replication Gateway pairs that you have chosen. You can apply the selected volume type and available zone either to all the virtual machines or customize for each.

Table 2-50 Volume type options available for Google Cloud Platform

Options	Description
Volume Type	Select the volume type for the disks. You can apply the selections either to all the virtual machines or customize for each. Note: The Regional Volume Type is not applicable for Boot disk. Even though the Regional Volume Type is selected in Apply All action, this volume type is not applicable for Boot disk.
IOPS (This field is enabled if Extreme Persistent Disk option is selected)	Enter the IOPS required if the volume type is Extreme Persistent Disk. Note: Refer to GCP documentation for more information on IOPS permitted for specific volume type and size.
Encryption Key	Select the Encryption Key, if you want to create encrypted volume in GCP. Before selecting key, ensure that key has all the required permissions. Note: Refer to Google Cloud Platform documentation troubleshooting section for the actual command in Resiliency Platform Product documentation. All the disks in Google Cloud Platform are already encrypted by Google Managed Key. Hence, you can choose to create and manage your own keys. These keys are also applicable for all or selected disks of a workload.
Zone	The zones which are listed are based on the Replication Gateway pairs that you have chosen. You can apply the selected volume type and available zone either to all the virtual machines or customize for each.
Replica Zone	This field is enabled when Regional disks types are selected.

Customize panel for Google Cloud Platform

This panel is displayed when you are configuring your virtual machines for recovery to Google Cloud Platform.

Table 2-51 Select Attributes panel

Options	Description
Use same firewall for rehearsal operation similar to the one in production checkbox	<p>Checkbox is selected by default.</p> <p>Machine type and Target VM values are already present.</p> <p>If the checkbox is uncheck, provide the following details:</p> <ul style="list-style-type: none"> ■ Network tag or Rehearsal Network Tag
Target VM Name	Value is already present. It can be changed later.
Machine Type dropdown	Select the value.
Network tag or Rehearsal Network Tag	These values are already present when the Use same firewall for rehearsal operation similar to the one in production checkbox is selected.

Select NIC panel

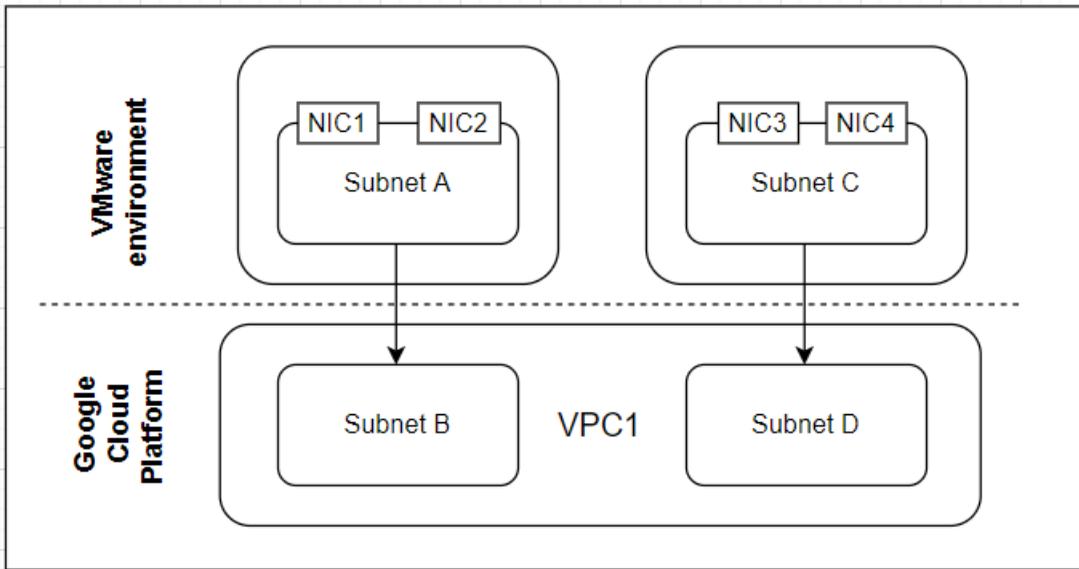
While configuring your virtual machines in Goggle Cloud Platform, each network interface of an instance must be attached to different VPC and also each network interface must belong to a different subnet. If network mapping conflicts to above requirement, then you must select **subnet** from dropdown and within the subnet select the primary NIC from the dropdown on **Select Primary NIC** panel and click **Next**.

Below is the example of network mapping which conflicts with Google Cloud Platform requirement:

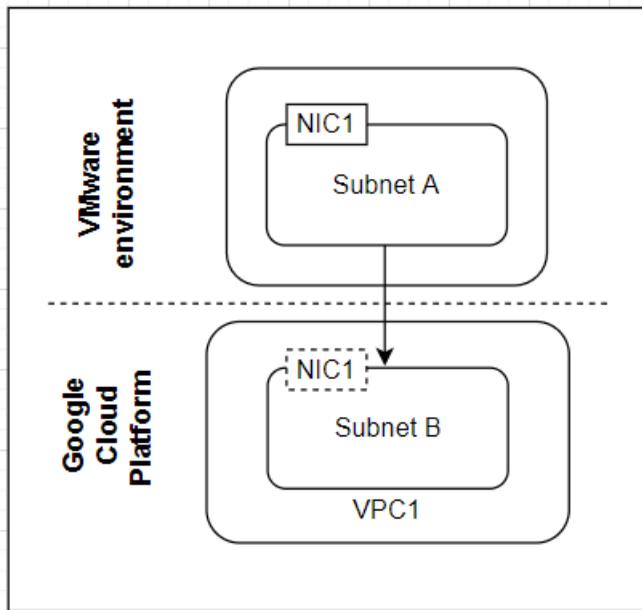
If you create network mapping from VMware or Hyper-V to Google Cloud Platform, say subnet A to B and subnet C to D and there are two NICs which exist in subnet A and two NICs exists in subnet C. On Google Cloud Platform, subnet B and subnet D belongs to same VPC say VPC1.

Below is the diagrammatical representation of the configuration explained:

Figure 2-8 Diagrammatical representation of the network configuration:



So considering the above configuration, the mapping in Google Cloud Platform should have each network interface attached to different VPC and so for different subnet there must be one NIC created out of above four NICs. Hence, as described in the below diagram if you select a subnet say Subnet A and then select a primary NIC say NIC1 from subnet A using dropdown in **Select Primary NIC** panel, after migrating the virtual machines the NIC1 is created on Subnet B in GCP environment.

Figure 2-9 Example of the configuration

Network customization options

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

To customize the static IP of Windows guest virtual machines in the VMware environment following are the two options:

- Use global user credentials for IP customization of Windows virtual machines. This option uses the Windows global user credentials. These credentials must be configured in advance.
- Install IP customization service on Windows virtual machines. After you finish installing the IP customization service, ensure that the following settings are disabled:
 - User Account Control: Admin Approval mode for the Build-in Administrator account
 - User Account Control: Run all administrators in Admin Approval mode. This option installs a service in the virtual machine to assist in IP customization. This service does not have any other functionality and does not require any

inbound or any outbound communication. The Resiliency Platform does not store the credentials which are provided to install the IP customization service. This option does not enforce authentication domain reachability on target data center unlike the global user option's requirement. This option is only applicable for VMware environment with third-party replication and Resiliency Platform Data Mover with VAIO framework.

You can do the following in this panel:

- Choose between the following two options for IP customization of Windows Virtual Machines
 - Use global user credentials for IP customization of Windows virtual machines.
 - Install IP customization service on Windows virtual machines.
- Manage PTR records
 - For Windows DNS and for Linux Bind, if you want Resiliency Platform to customize DNS settings, then DNS records should not exist at the target data center.
 - For Windows DNS, if you configure a user in Resiliency Platform for DNS customization, then that user should also have rights to update DNS records added by any other user on the DNS.
- Choose to continue with DR operations even if DNS updates fail.

You can customize the IPs for Production and Rehearsal networks. Customizing the IPs of a virtual machine overrides the default IP settings when the virtual machine starts at the target data center.

If the subnet mapping is already done, and the mapped subnets are of equal mask, then the computation of projected static IP is done based on the subnet mappings by Resiliency Platform. This projected IP address for the target data center can be edited. If the mapped subnets are of unequal masks, then you need to enter the IP address manually.

If the subnet mapping is not done, then you can either select a subnet from the drop-down list and apply it to all the target IP addresses or you can select a separate subnet for each target IP address. You also need to enter the IP address for the target data center. Since IPv6 network support is enabled and subnets can be created using IPv6 address, you can see the IPv6 subnets in the drop-down. Either you can apply the IPv6 address to all the target IP address or you can select a separate subnet for each target IP address.

You can choose to continue with the DR operation if the IP customization fails. Note that this is possible only if the virtual machines have static IPs.

To customize the static IP of Windows guest virtual machines in the VMware environment, Resiliency Platform requires the user name and password to log on to the Windows virtual machines. To configure this user name and password go to **Product Settings > User Management > Windows Global User**.

See [“Configuring Windows global user”](#) on page 482.

If the target data center is in cloud, then ensure that the IPs used for network customization are not already in use on the cloud.

If you choose to apply DNS customization, then you can add a host name to IP mapping of the DNS.

About manual intervention

Resiliency Platform lets you pause the Migrate and Recover operation at certain predefined stages. This gives you time to perform any manual tasks on the assets if required. On completion of your tasks, you can resume the operation from **Activities** or **Recent Activities** menu.

When you are configuring a resiliency group for disaster recovery (DR), one of the panels lets you select the pause or the manual intervention point. Similarly when you are configuring a Virtual Business Service (VBS) for DR you can choose the manual intervention points.

Following are the predefined points while configuring a resiliency group:

- After stopping the assets on the source data center.
- Before registering the assets on the target (recovery) data center.

While configuring a VBS, the predefined points are during stop of a tier and start of a tier.

- Stop of a tier: Before stopping and after starting the workloads.
- Start of a tier: Before stopping and after starting the workloads.

Enabling manual intervention is an optional step. This panel is displayed when you are configuring the resiliency group for the following scenarios:

- Recovery to on-premises data center using Resiliency Platform Data Mover.
- Recovery to on-premises data center using any 3rd party replication technology.
- Recovery to any cloud data center.

Note: The Migrate and Recover operations remain in paused state till you resume the operations.

You can view the time spent for manual intervention in the following reports:

- Activity Distribution History
- Recovery Activity History by RG
- Recovery Activity History by VBS

See [“Managing a running activity”](#) on page 582.

Advanced features

Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.

- See [“About virtual business services”](#) on page 526. Managing virtual business services
- See [“About resiliency groups with assets”](#) on page 528. Managing resiliency plans
- See [“About evacuation plan”](#) on page 530. About evacuation plan

About virtual business services

For a business service to work properly, it is important that all of its tiers and components are up and working together. From a business continuity point of view, it is important to not just ensure that individual tiers are up and running but also the entire business service.

A virtual business service (VBS) is a logical collection of resiliency groups that function together to perform a particular business service. A VBS enables easy management of multi-tier business services. For example, you can group a web server resiliency group, a database resiliency group, and a payroll business logic resiliency group into a VBS called `payroll`. You can start, stop, monitor, manage, or recover that VBS as a single entity.

Note: If a VBS consists of resiliency groups that are in maintenance mode, then you cannot perform any operations on the VBS.

Understanding tiers

Within a VBS, resiliency groups are arranged in tiers. Tiers represent the logical dependencies between the resiliency groups and determine the relative order in which the resiliency groups start and stop. For example, the database resiliency group must start before the application server resiliency group and the web server resiliency group, so the database resiliency group must go in the lowest tier. The

application server resiliency group must start after the database resiliency group, so it goes in the next tier. The web server resiliency group must start last, so it goes into the top tier. Later, if you add a resiliency group to the VBS, you can manage it as part of the IT service by placing it in the appropriate tier.

Various configuration states of virtual business services

It may happen that during upgrade or configuring VBS, the workflow may stuck at some particular state where manual reconfiguration is required.

Table 2-52 VBS configuration states and resolution

Configuration states	Virtual business service may lead into the state on following events	Resolution
Configuring	<ul style="list-style-type: none"> ■ Create virtual business service ■ Edit virtual business service ■ Edit resiliency group associated with virtual business service ■ Delete resiliency group associated with virtual business service 	N/A
Configuration failed	<ul style="list-style-type: none"> ■ Create virtual business service fails ■ Edit virtual business service fails ■ Virtual business service reconfiguration triggered by following event fails: <ul style="list-style-type: none"> ■ Edit resiliency group associated with virtual business service ■ Delete resiliency group associated with virtual business service 	Edit the virtual business service manually.

Table 2-52 VBS configuration states and resolution (*continued*)

Configuration states	Virtual business service may lead into the state on following events	Resolution
Unconfiguring	Delete virtual business service	N/A
Unconfiguration failed	Delete virtual business service fails	Retry deleting the virtual business service.
Refreshing	Refresh virtual business service	N/A
Reconfiguration pending	<p>Virtual Business Service is currently in Configuring/Refreshing state and at the same time Reconfiguration of virtual business service is triggered by event:</p> <ul style="list-style-type: none"> ■ Edit resiliency group associated with virtual business service ■ Delete resiliency group associated with virtual business service 	<p>N/A</p> <p>Note: This is an intermediate state and should get resolved automatically.</p>

See [“About virtual business services”](#) on page 526.

About resiliency groups with assets

Resiliency groups are the unit of management and control in Veritas Resiliency Platform. After assets are added to Resiliency Platform, you organize related assets into a resiliency group that you can protect and manage as a single entity. A resiliency group can have only physical machines or only virtual machines, a mix of physical and virtual machines is not supported. Similarly it can contain either all applications or all InfoScale applications but not a mix of both.

For example, you can organize several physical or virtual machines into a resiliency group, and name it `VM_Finance`. When you perform an operation on the `VM_Finance` resiliency group using the Resiliency Platform console, the operation is performed on all the assets that belong to the resiliency group. For example if you run the Start operation on the resiliency group, all the assets (physical machines or virtual

machines) that belong to the resiliency group start booting. Or if you perform any of the disaster recovery operations such as Migrate on the resiliency group, all the assets within the group are migrated to the selected target data center.

Ensure that the following prerequisites are met while creating a resiliency group for remote recovery:

- Ensure that all the virtual machines that are to be grouped in a single resiliency group belong to a single hypervisor or virtualization server (if not clustered), or to a single cluster.
- Encryption and compression are disabled on Hyper-V servers.
- Refresh the virtualization server (vCenter server or Hyper-V server), the host, and the cloud discovery before proceeding to create the resiliency group.
- If the recovery is on Azure, then ensure that the virtual machine names should always start and end with alphanumeric characters. The name can contain periods (.), hyphens(-), or underscores(_) in the middle.
- If you are using third-party replication, ensure that the assets consume storage from the same consistency groups. E.g. EMC SRDF device group, NetApp Volume, 3PAR replication group, and so on.

The operations available for a resiliency group depend on how it is configured. While configuring a resiliency group, you need to select a service objective. If you select a service object that supports remote recovery, then you can perform disaster recovery operations such as Migrate and Take over on the resiliency group.

Optionally you can use a service objective that only monitors the assets or the applications and provides only basic operation capabilities like start and stop operations and no remote recovery operations. Using the Copy service objective, you can recover the virtual machines from NetBackup generated backup images to the target data center.

See [“Configuring a resiliency group for basic monitoring”](#) on page 529.

Configuring a resiliency group for basic monitoring

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

There are two types of pre-activated service objectives:

- Monitor assets - provides only monitoring, start, and stop operations
- Recover hosts - provides recovery operations as well as the start and stop operations

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

To manage assets for basic monitoring

1 Prerequisites

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

2 Navigate



Assets (navigation pane) **Unmanaged** tab > **Manage & Monitor Assets**

You can also launch the wizard from the **Overview** tab.

3 Select the assets:

- Select **Host** as the asset type, select the data center, type, and other filters as needed to display a list of assets.
- Drag and drop the selected assets to **Selected Instances**.

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

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

6 Supply a name for the resiliency group.

7 Verify that the new resiliency group is added to the **Resiliency Group(s)** tab.

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

Note: If the instances are created from BYOS image or there is licensing issues with instances, then Resiliency Platform operations may fail.

See [“About resiliency groups with assets”](#) on page 528.

About evacuation plan

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

Using the evacuation plan template you can define the sequence in which the virtual business services (VBS) should be migrated from the production data center to the recovery data center. Resiliency groups that do not belong to any VBSs, are appended at the end of the evacuation plan workflow after the VBS. If there are large number of VBSs then up to 5 VBSs within a priority group are migrated in

parallel to the recovery data center. Similarly, if there are large number of resiliency groups, up to 10 resiliency groups are migrated in parallel.

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

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

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

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

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

For a VBS or a resiliency group to successfully evacuate to the target data center, it should meet the following criteria:

- VBS or resiliency group that belong to the evacuation plan must be configured for disaster recovery.
- VBS can contain resiliency groups some of which are configured for disaster recovery and some using the service objective with data availability as Copy.
- Resiliency group must belong to only one VBS.

When you generate a plan, an appropriate warning is shown listing the assets that are excluded from the plan.

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

- Evacuate
- Rehearse evacuation
- Cleanup evacuation rehearsal
- Regenerate

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

- VBSs are added, modified, or deleted.
- Resiliency groups are added and configured for disaster recovery.
- Resiliency groups which were configured for disaster recovery are deleted.

- Existing resiliency group is configured for disaster recovery.

No action is required in the following scenarios:

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

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

Perform remote recovery operations

Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.

- See [“Performing the rehearsal operation for virtual machines”](#) on page 532. Performing the rehearsal operation for virtual machines
- See [“Performing cleanup rehearsal for virtual machines”](#) on page 535. Performing cleanup rehearsal for virtual machines
- See [“Migrating a resiliency group”](#) on page 536. Migrating a resiliency group
- See [“Performing the resync operation for virtual machines”](#) on page 537. Recovering resiliency group of virtual machines
- See [“Recovering a resiliency group using replication-based recovery”](#) on page 539. Performing the resync operation for virtual machines

Performing the rehearsal operation for virtual machines

Use the **Rehearsal** option on the Resiliency Platform console to ensure the disaster recovery readiness of the assets in your protected resiliency groups.

From version 3.5, Rehearsal operation can be performed on Azure or (Azure Stack) data center when it is deployed as source as well as a target data center.

For recovery on AWS cloud:

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

during Rehearsal is significantly reduced. Which reduces the overall time taken to complete the operation.

Note: This setting is specific to recovery of virtual machines from VMware to VMware if Resiliency Platform Data Mover is used and recovery of physical machines to VMware data center:

While performing the rehearsal operation, DRS automation level for the target Replication Gateway should be set to manual or it should be disabled.

To perform the rehearsal operation

1 Prerequisites

See [“Prerequisites for rehearsal operation for virtual machines”](#) on page 534.

2 Navigate



Assets (navigation pane) > **Resiliency Group(s)** tab

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

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

Rehearsal operation on virtual machines for CDP

The rehearsal option on the Resiliency Platform console is used to ensure the disaster recovery readiness of the assets in your protected resiliency groups. You can perform the rehearsal operation for the resiliency groups where CDP is enabled.

To perform rehearsal operation for virtual machines for CDP

Navigate

1



Assets (navigation pane) > **Resiliency Group(s)** tab

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

3 Select the target data center and then select the recovery points to perform the rehearsal operation.

4 Click **Next**.

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

See [“Performing cleanup rehearsal for virtual machines”](#) on page 535.

Prerequisites for rehearsal operation for virtual machines

Before you run the rehearsal operation for a resiliency group, ensure that you have met the following prerequisites:

- For VMware virtual machines, ensure that the datastores have enough free space for the swap files for the on-premises virtual machines and the virtual machines created by the rehearsal operation on the recovery data center. The size of the swap files is same as that of the virtual machine memory size.
- For VMware virtual machines, ensure that the mapping of all the required port groups across the data centers is complete.
For Hyper-V virtual machines, ensure that the mapping of all the required virtual switches across the data centers is complete.
See [“Creating network pairs between source and target data centers”](#) on page 514.
- Each type of replication has prerequisites and limitations for the rehearsal operation.
- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs. This is required if the recovery data center is AWS, Azure cloud.
- If the recovery data center is in AWS, then configure a rehearsal subnet in the cloud. The rehearsal and production subnet should be in the same VPC.
- If the recovery data center is in cloud, then you need to set the SAN policy to either OnlineAll or OfflineShared based on whether you have shared or non-shared disks. For more details refer to [Microsoft Documentation](#).
- If the status of the virtual machine on the recovery data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.
- When the **Use Same WWN** option is checked, Resiliency Platform needs to do additional storage operations to map or unmap the LUNs, swap WWNS of LUNs during the workflow and that requires extra rescans and validations. These operations are not required when the option **Use Same WWN** is disabled. This option can be set only for when the DR operation is invoked from Resiliency Platform. It can be reset once the DR activity is complete.

Note: This setting is specific to recovery of virtual machines from VMware to VMware if Resiliency Platform DataMover is used and recovery of physical machines to VMware data center:

While performing the rehearsal operation, DRS automation level for the target Replication Gateway should be set to manual or it should be disabled.

See [“Performing the rehearsal operation for virtual machines”](#) on page 532.

Performing cleanup rehearsal for virtual machines

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

Note: Any snapshots of the cloud volumes that are taken external to Veritas Resiliency Platform may cause failure in rehearsal cleanup. This is applicable only for Orange Recovery Engine.

This setting is specific to recovery of virtual machines from VMware to VMware if Resiliency Platform data Mover is used and recovery of physical machines to VMware data center: While performing the rehearsal operation, DRS automation level for the target Replication Gateway should be set to manual or it should be disabled.

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

- A separate copy of virtual machine when you use Hyper-V Replica for data replication.
- A new registered virtual machine that has its virtual machine data files (VHDX) residing on snapshot LUNs when array-based replication (for example, EMC SRDF) is used for data replication.

Using NetBackup

When your assets are configured for remote recovery using a service objective where the data availability mode is Copy, then during the rehearsal operation virtual machines are created on the recovery data center with the selected backup image. These virtual machines and the data are deleted during the cleanup operation.

To perform cleanup rehearsal

- 1 Navigate to **Assets** (navigation pane) > **Resiliency Group(s)** tab.
- 2 Double-click the resiliency group to view the details page. Click **Cleanup Rehearsal**.
- 3 Select the target data center, and click **Next**.

See [“Performing the rehearsal operation for virtual machines”](#) on page 532.

Migrating a resiliency group

Migration refers to a planned activity involving graceful shutdown of physical and virtual machines at the source data center and starting them at the target data center. In this process, replication ensures that consistent data of the assets is made available at the target data center which could be the on-premises or cloud data center. In Veritas Resiliency Platform, the migration of assets is achieved by grouping them in a resiliency group, configuring disaster recovery for the resiliency group, and thereafter performing the migrate operation on this resiliency group.

Consider the following:

- If you perform the recover operation, then you must perform the Resync operation before you migrate back to the production data center.
- If the **Enable reverse replication** option is not selected, then before migrating the virtual machines to the target data center and after migrating back to the source data center, you need to perform the Resync operation. See [“Performing the resync operation for virtual machines”](#) on page 537.
- If the recovery data center is Azure cloud, then after you migrate from the cloud data center to the on-premises data center, you need to refresh Azure cloud to rediscover the cloud-based objects.
- During the migrate operation, virtual machines on the source data center are gracefully shut down. If you manually shut down the virtual machine before performing the migrate operation, and if the shut down was not graceful, then the migrate operation may fail. This is applicable when the replication technology is Resiliency Platform Data Mover and the target data center is in cloud.
- When you upgrade from an earlier version to version 3.2 or later, then after performing the migrate operation, a risk is raised. This risk is regarding the changes in NIC configuration when you migrate to any cloud data center. Suppress this risk while the resiliency group is online on cloud. Migrate back to the on-premises data center and then edit the resiliency group to fix the NIC configuration.

Prerequisites

- Ensure that the Data Mover connection status is **Connected**, Data State is **Consistent**, and Replication State is **Active**.
- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs. This is required if the recovery data center is AWS, Azure or Google Cloud Platform.
- For VMware virtual machines, ensure that the network mapping of all the required port groups, or subnets across the data centers is complete.

For Hyper-V virtual machines, ensure that the network mapping of all the required virtual switches across the data centers is complete.

See “[Creating network pairs between source and target data centers](#)” on page 514.

- If the recovery data center is in AWS, Azure or Google Cloud Platform then ensure that the network mapping of all the required subnets across the data centers is complete.
- If the recovery data center is in cloud, then you need to set the SAN policy to either OnlineAll or OfflineShared based on whether you have shared or non-shared disks. For more details refer to [Microsoft Documentation](#).
- If the status of the virtual machine on the recovery data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.
- For the replication technology HPE 3PAR Remote Copy, ensure that for VMware virtual machines the `config.vpxd.filter.hostRescanFilter` value is set to false.

To migrate a resiliency group

1 Navigate



Assets (navigation pane) > **Resiliency Group(s)** tab

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

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

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

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

For more information on troubleshooting specific scenarios,

Performing the resync operation for virtual machines

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

Since the production data center is not working, the data replication between the two sites does not happen. After the production site is back up and running, you

need to prepare the production site for the next failover or for a migration operation. This preparation includes cleaning up any residue and resuming the replication from the recovery to the production site.

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

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

Consider the following if you have configured your assets for recovery to vCloud Director without adding the VMware vCenter server or Hyper-V server:

You need to perform the resync operation after Migrate or Recover. In the **Activities** panel, if the workflow is in **Paused** state for "Resync Replication" subtask, then you need to manually start the physical and virtual machines. Ensure that the physical and virtual machines boot from PXE OS of the replication gateway that is configured as PXE server. You can verify the virtual machines boot progress from vCenter or Hyper-V console by checking the virtual machine console. When the boot is complete, click **Resume** so that the workflow proceeds with synchronizing data from the disks on recovery datacenter to those on the on-premises data center. After Resync operation is complete, do not shut down the virtual machines, otherwise the subsequent Migrate or Recover operations fail.

Performing the resync operation

1 Prerequisites

- If the target (recovery) data center is on-premises, and the last performed operation was recover, then you may need to restart the Hyper-V server or the ESX server. Although the Resync operation cleans up any residue on the source data center before resuming replication, there could be some residues that can be cleaned up only by restarting the hypervisors.
- For the replication technology HPE 3PAR Remote Copy, ensure that for VMware virtual machines the `config.vpxd.filter.hostRescanFilter` value is set to false.
- If the status of the virtual machine on the recovery data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.

Prerequisites for resync operation of physical workloads

- The resiliency group must be configured for DR and migrated to the recovery datacenter.

- Ensure that PXE boot server is configured on the on-premises Replication Gateway.
 - Ensure to set network boot (PXE boot) as the first boot priority in the system BIOS.
 - If you are not using 3rd party DHCP server, then configure DHCP server on the PXE Boot server that is configured on the on-premises Replication Gateway.
- 2 Navigate to **Assets** (navigation pane) > **Resiliency Group(s)** tab
 - 3 Double-click the resiliency group to view the details page. Click **Resync**.
 - 4 In the **Resync** panel, select the production data center name from the drop-down list, and click **Next**.

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

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

After completing recover from cloud and resync, the resiliency group details page shows entries for deleted or unavailable virtual machines on cloud data center. To remove these stale entries, after resync is complete, edit the resiliency group with **Edit Configuration** intent. You may submit the wizard without making any changes.

Recovering a resiliency group using replication-based recovery

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

Perform the resync operation after successful completion of recover operation.

If you are recovering to vCloud Director data center, without adding Hyper-V Server or vCenter Server, then recover operation from cloud to production (on-premises) data center is not supported.

Prerequisites

- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs. This is required if the recovery data center is AWS, Azure cloud.
- For VMware virtual machines, ensure that the network mapping of all the required port groups, or subnets across the data centers is complete.
For Hyper-V virtual machines, ensure that the network mapping of all the required virtual switches across the data centers is complete.
See [“Creating network pairs between source and target data centers”](#) on page 514.
- If the source data center is in AWS, then ensure that the network mapping of all the required subnets between the source and target data center is complete.
- There should not be any resources on Azure having the same name or substring of name as that of the virtual machine display name or FQHN name on on-premises data center.
- If the source data center is in cloud, then you need to set the SAN policy to either OnlineAll or OfflineShared based on whether you have shared or non-shared disks. For more details refer to [Microsoft Documentation](#).
- If the status of the virtual machine on the source data center is not correctly displayed, then you need to refresh the cloud discovery or the virtualization server discovery.
- For the replication technology HPE 3PAR Remote Copy, ensure that for VMware virtual machines the `config.vpxd.filter.hostRescanFilter` value is set to false.

When you upgrade from an earlier version to version 10.0 or later, then after performing the recover operation with replicated data, a risk is raised. This risk is regarding the changes in NIC configuration when you recover to any cloud data center. Suppress this risk while the resiliency group is online on cloud. Migrate back to the on-premises data center and then edit the resiliency group to fix the NIC configuration.

To perform recover operation on virtual machines

1 Navigate



Assets (navigation pane) > **Resiliency Group(s)** tab

- 2 Double-click the resiliency group to view the details page. Click **Recover**.
- 3 Do the following:

- Select the target data center.
- If there is an outage on the source data center, select the **Confirm outage of assets** check box.
- During the recover operation, if Resiliency Platform detects a probability of data loss, you have the option to abort the recover operation to avoid any data loss. Select the check box if you want to abort the operation in such a situation.
- Click **Continue** for warnings.

4 Click **Submit**.

To perform recover operation on resiliency group configured with CDP enabled

1 Navigate



Assets (navigation pane) > **Resiliency Group(s)** tab

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

3 To select the recovery point for CDP, do the following:

4 On the **Select Recovery points** panel, you can select the date, time, and range to list the recovery points with the latest data against the assets.

- Select the target data center.
- If there is an outage on the source data center, select the **Confirm outage of assets** check box.
- Do not select the **Abort recover if these subtasks fail** check box and click **Next**.
- Select the recovery points. To choose from a specific time range, select the date and the time range. Then select the hours or minutes since the start time. Click **Search**.

5 Click **Submit**.

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

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

For more information on troubleshooting specific scenarios,

Monitor assets

You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.

- See [“About risks”](#) on page 542.About risks
- See [“About reports”](#) on page 576.About reports
- See [“Managing a running activity”](#) on page 582.Managing activities

About risks

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

Risk Insight detects the changes to the state and configuration of your protected assets. It identifies if there is a risk to the recoverability or continuity of your protected assets. A periodic or schedule scan will generate the risk on the Resiliency Platform components. A periodic scan of every 30 minutes is run by the scheduler.

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

Risk insight generates two types of reports:

- **Current risk reports:** Provides the summary and detail information about all the current risks in your data center.
- **Historical risk reports:** Provides a summary and a detailed analysis of information about the risks in your environment during the specified period.

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

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

Table 2-53 Risk types

Risk category	Description
Recoverability	Risks that may impact the ability to recover and run the application on the recovery site.

Table 2-53 Risk types (*continued*)

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

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

Table 2-54 Risk types

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

From 3.2 onwards, you can probe and suppress a risk based on which component of Resiliency Platform the risk has occurred.

Probing a risk is way of evaluating a risk to check whether the risk is eliminated or still exists at the Resiliency Platform component. You cannot probe all the risks in the risk view. The risks which cannot be probed have details about the risk resolution. You can perform the given steps and probe the risk again. You have to wait for some time the to see the apply the changes.

For example, an IMS is disconnected due to network issues. The risk is then raised by the Resiliency Platform. When you probe this risk, details of this risk display what steps can be done to resolve this risk. There are some risks which cannot be probed. For those risks, you have to fix the risk and the probe it accordingly.

When you probe a risk, the details panel also displays a link to the proposed resolution for the risk. On clicking the link, the risk pop up is closed and you are redirected to the respective resolution page.

You can avoid a risk using suppress option on the risk view. You can suppress a risk for specific time; for few minutes or hours but cannot suppress a risk for indefinite period. If a risk is active after the suppress period is over, then it will appear under the Active Risks view. You can probe that risk or you can again suppress it for some time. Suppress option is available for all the risk in the Resiliency Platform. Suppressed risks will be seen in Risk view > Suppressed Risks.

Note: Risks are not generated for resiliency groups that are in maintenance mode. If a risk was raised before the resiliency group was placed in maintenance mode and that risk persists after exiting the mode, then the risk is shown.

See [“Predefined risks in Resiliency Platform”](#) on page 544.

Predefined risks in Resiliency Platform

[Table 1-72](#) lists the predefined risks available in Resiliency Platform. These risks are reflected in the current risk report and the historical risk report.

Table 2-55 Predefined risks

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
vCenter Password Incorrect	Checks if vCenter password is incorrect	15 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop operations ■ On secondary site: migrate or recover operations 	In case of a password change, resolve the password issue and refresh the vCenter configuration
VM tools not installed	Checks if VM Tools are not Installed. It may affect IP Customization and VM Shutdown	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Stop 	<ul style="list-style-type: none"> ■ In case of VMWare, install VMWare Tools ■ In case of Hyper-V, install Hyper-V Integration Tools
Snapshot reverted on Virtual Machine	Checks if snapshot has been reverted on virtual machine	5 minutes	Error	Resiliency Platform Data Mover replication	Perform the Resync operation on the resiliency group.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Resiliency Platform Data Mover daemon crashed	Resiliency Platform Data Mover filter is not able to connect to its counterpart in ESX. The replication process has stopped and is at risk	5 minutes	Error	Resiliency Platform Data Mover replication	<ul style="list-style-type: none"> ■ To continue the replication, you can move (VMotion) the virtual machine to a different ESX node in the cluster. ■ Troubleshoot the issue with this ESX node or raise a support case with Veritas.
DataMover virtual machine in no-op mode	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Veritas Replication policy has been detached	Veritas Replication policy has been detached from the disk associated with virtual machine.	5 minutes	Error	Migrate	Perform Resync operation on the affected resiliency group.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Asset disk configuration changed	Checks if disk configuration of any of the assets in the resiliency group has changed.	30 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal 	<p>Refresh the respective hosts, vCenter servers or Hyper-V servers and the cloud discovery. After refresh, probe the risk.</p> <p>After performing the above mentioned step even if the risk still exists, edit the resiliency group to first remove the impacted virtual machine from the resiliency group and then add it back to the resiliency group.</p>
Asset NIC configuration changed	Checks if NIC configuration of any of the assets in the resiliency group has changed.	30 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Resync 	<p>If the resilience group is online on the target data center, then either revert the NIC changes done on the virtual machines or suppress the risk to be able to migrate the assets back to the source data center. If the resiliency group is online on source data center, edit the resiliency group with Edit Configuration or Customize Network option to update the NIC configuration.</p>

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Invalid NIC Configuration	One or more NICs on the host are not configured properly.	Real time, while creating resiliency group	Error	Create resiliency group	Ensure that the keys NAME, DEVICE and HWADDR have appropriate values as per the details of each NIC in its configuration file.
Global user deleted	Checks if there are no global users. In this case, the user will not be able to customize the IP for Windows machines in VMware environment	Real time	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Edit the resiliency group or add a Global user
Failure to validate Windows Global User credentials for IP customization	<p>This risk is raised if:</p> <ul style="list-style-type: none"> ■ Windows Global User is not configured. ■ Windows Global User does not have appropriate credentials. ■ Virtual machine is offline while configuring resiliency group for disaster recovery. 	After the resiliency group is configured for disaster recovery	Warning	<ul style="list-style-type: none"> ■ Rehearsal ■ Migrate ■ Recover 	Add Windows Global Users with appropriate credentials. Edit the resiliency group using the Network Customization option to resolve the risk.
Missing heartbeat from Resiliency Manager	Checks for heartbeat failure from a Resiliency Manager	5 minutes	Error	All	Fix the Resiliency Manager connectivity issue
Infrastructure Management Server disconnected	Check for Infrastructure Management Server(IMS) to Resiliency Manager(RM) connection state	1 minute	Error	All	Check IMS reachability Try to reconnect IMS
Storage Discovery Host down	Checks if the discovery daemon is down on the storage discovery host	15 minutes	Error	Migrate	Resolve the discovery daemon issue

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
DNS removed	Checks if DNS is removed from the resiliency group where DNS customization is enabled	real time	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Edit the Resiliency Group and disable DNS customization
IOTap driver not configured	Checks if the IOTap driver is not configured	2 hours	Error	None	Configure the IOTap driver This risk is removed when the workload is configured for disaster recovery.
VMware Discovery Host Down	Checks if the discovery daemon is down on the VMware Discovery Host	15 minutes	Error	Migrate	Resolve the discovery daemon issue
VM restart is pending	Checks if the virtual machine has not been restarted after add host operation	2 hours	Error	Create resiliency group	Restart the virtual machine after add host operation
New virtual machine added to replication storage	Checks if a virtual machine that is added to a Veritas Replication Set on a primary site, is not a part of the resiliency group	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Rehearsal 	Add the virtual machine to the resiliency group
Replication lag exceeding RPO	Checks if the replication lag exceeds the thresholds defined for the resiliency group. This risk affects the SLA for the services running on your production data center	5 minutes	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Check if the replication lag exceeds the RPO that is defined in the Service Objective
Replication state broken/critical	Checks if the replication is not working or is in a critical condition for each resiliency group	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Contact the enclosure vendor. In case of Resiliency Platform Data Mover, or raise a support case with Veritas

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Remote mount point already mounted	Checks if the mount point is not available for mounting on target site for any of the following reasons: <ul style="list-style-type: none"> ■ Mount point is already mounted ■ Mount point is being used by other assets 	<ul style="list-style-type: none"> ■ Native (ext3, ext4, NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours 	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Unmount the mount point that is already mounted or is being used by other assets Risk gets resolved after 30 minutes if a successful cleanup rehearsal, migrate, or recover operation performed and VMware vCenter gets refreshed within 30 minutes.
Disk utilization critical	Checks if at least 80% of the disk capacity is being utilized. The risk is generated for all the resiliency groups associated with that particular file system	<ul style="list-style-type: none"> ■ Native (ext3, ext4, NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours 	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Rehearsal 	Delete or move some files or uninstall some non-critical applications to free up some disk space
ESX not reachable	Checks if the ESX server is in a disconnected state	5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop operations ■ On secondary site: migrate or recover operations 	Resolve the ESX server connection issue

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
vCenter Server not reachable	Checks if the virtualization server is unreachable or if the password for the virtualization server has changed	5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop operations ■ On secondary site: migrate or recover operations 	Resolve the virtualization server connection issue In case of a password change, resolve the password issue
vCenterDown	1.The vCenter server is down or unreachable. 2. The vCenter server is down. Unable to establish secure connection with vCenter server as the SSL/TLS handshake has failed.	15 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop, migrate, rehearsal, local recover, resync operations. ■ On secondary site: migrate, resync, recover, cleanup rehearse operations. 	For pt 1. Check for any one of the following issues and resolve: <ol style="list-style-type: none"> 1 vCenter server is down, vSphere service is not working, or vCenter server port has been changed after vCenter server configuration. 2 If port has been changed after configuration then, perform edit vCenter server operation to resolve the issue. For pt 2. Install valid CA certificates of vCenter server in the Resiliency Platform and refresh the vCenter server configuration

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Insufficient compute resources on failover target	Checks if there are insufficient CPU resources on failover target in a virtual environment	6 hours	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Recover 	Reduce the number of CPUs assigned to the virtual machines on the primary site to match the available CPU resources on failover target
Host not added on recovery data center	Checks if the host is not added to the IMS on the recovery data center	30 minutes	Error	Migrate	Check the following and fix: <ul style="list-style-type: none"> ■ Host is up on recovery data center ■ Host is accessible from recovery datacenter IMS ■ Time is synchronized between host and recovery datacenter IMS
NetBackup Notification channel disconnected	Checks for NetBackup Notification channel connection state	5 minutes	Error	Recover	Check if the NetBackup Notification channel is added to the NetBackup primary server If the risk resolution or description indicates an SSL/TLS verification error. Refer this troubleshooting guide.
Backup image violates the defined RPO	Checks if the backup image violates the defined RPO	30 minutes	Warning	No operation	<ul style="list-style-type: none"> ■ Check the connection state of NetBackup Notification channel ■ Check for issues due to which backup images are not available

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
NetBackup primary server disconnected	Checks if NetBackup primary server is disconnected or not reachable	5 minutes	Error	Recover	Check if IMS is added as an additional server to the NetBackup primary server
NetBackup Recovery Host decommissioned	Check if NetBackup Recovery Host is disconnected or not reachable.	5 minutes	Error	recover	<ul style="list-style-type: none"> ■ Edit the resiliency group and choose different recovery host. ■ Try to connect the same recovery host.
Assets do not have copy policy	Checks if the assets do not have a copy policy	3 hours	Warning	No operation	Set up copy policy and then refresh the NetBackup primary server
Target replication is not configured	Checks if the target replication is not configured	3 hours	Warning	No operation	Configure target replication and then refresh the NetBackup primary server
Disabled NetBackup Policy	Checks if NetBackup policy associated with the virtual machine is disabled	3 hours	Warning	No operation	Fix the disabled policy

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Replication block tracking disk not found	Checks for the replication block tracking disk. If the replication block tracking disk is not found, then virtual machine does not get configured for remote recovery and the replication stops	30 minutes	Error	Migrate	Ensure that the RBT disk is attached to the virtual machine. After the risk gets resolved, perform reboot of VM then perform the resync operation to avoid disk corruption during migrate or migrate back. If you are not able to locate the RBT disk then perform following steps in the order listed: <ol style="list-style-type: none"> 1 Remove the virtual machine from resiliency group. 2 Add it again to a resiliency group to ensure that virtual machine is protected.
Members are manually deleted from network groups	Network group goes into faulted state when a member is manually removed. The risk is circulated to resiliency group	Immediate	Warning	Migrate, Rehearse	Edit the network group by adding the missing member and then edit the resiliency group details
Members deleted from network groups	Network group goes into faulted state when a discovered member gets deleted from IMS. The risk is circulated to resiliency group	5 minutes	Warning	Migrate, Rehearse	Edit the network group by adding the missing member and then edit the resiliency group details

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Virtual machine configuration not backed up	Unable to take a backup of virtual machine configuration file.	Immediate	Error	<ul style="list-style-type: none"> ■ Create resiliency group ■ Migrate ■ Rehearse 	Check the state of the IMS and its corresponding assets such as the hypervisors and vCenter servers. Perform edit resiliency group operation.
Unable to backup latest Virtual machine configuration	Unable to take a backup of the latest configuration file of the virtual machine.	Immediate	Warning	<ul style="list-style-type: none"> ■ Edit resiliency group ■ Migrate ■ Rehearse 	Check the state of the IMS and its corresponding assets such as the hypervisors and vCenter servers. Perform edit resiliency group operation.
Datastore for disk has changed to X, this datastore is not part of resiliency group	If virtual disk is moved to a non-compliant datastore. Applicable for 3rd party replication technology	5 to 15 minutes	Error	All operations except start and stop resiliency group	Edit the resiliency group or move the disk to a datastore which is part of the resiliency group.
Datastore for configuration file has changed to X, this datastore is not part of resiliency group. Previous datastore was Y.	If the virtual machine configuration file is moved to a non-compliant datastore. Applicable for 3rd party replication technology	5 to 15 minutes	Error	All operations except start and stop resiliency group	Edit the resiliency group or move the disk to a datastore which is part of the resiliency group.
Disk path has changed	Displayed when virtual machine snapshot is taken. Risk is resolved automatically after updating the blob.	5 to 15 minutes	Error	All operations	Risk is automatically resolved.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
New datastore added to the consistency group is not part of resiliency group	New datastore added to consistency group Applicable for 3rd party replication technology	6 hours	Error	<ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync 	Edit the resiliency group
Datastore removed from resiliency group	Datastore removed from consistency group Applicable for 3rd party replication technology	6 hours	Error	<ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync 	Edit the resiliency group
Veritas Replication VIB upgrade pending	Checks if the Veritas Replication VIB version on ESXi cluster has latest version installed.	6 hours	Error	None	Upgrade the Veritas Replication VIB to the latest version.
Veritas Replication VIB is in partial state.	Checks if the Veritas Replication VIB installation on ESXi cluster is in partial or unknown state.	6 hours	Error	<ul style="list-style-type: none"> ■ If the risk is on the target ESXi cluster then block the migrate and rehearsal operations. ■ If the risk is on the source ESXi cluster then block the resync operation. 	Perform Resolve and Verify operation on the ESXi cluster to fix the installation issues.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Insufficient privileges on vCenter server	Operations on the resiliency group may fail because of missing privileges on vCenter server data centers.	6 hours	Warning	One or more operations on resiliency group may fail because of missing privileges on vCenter server data center.	Ensure that appropriate privileges are configured on vCenter server data center before invoking any operation. Refer to the documentation for the required privileges.
Infrastructure Management Server data reporting disabled	Infrastructure Management Server cannot report data to Resiliency Manager due to version incompatibility	As soon as IMS connects to the Resiliency Manager after the Resiliency Manager upgrade	Error	All	Upgrade IMS to the latest version that is specified in the risk message
DRS Datastore Is Added Or Removed	New datastore is added to the cluster or is removed from the cluster	6 Hours	Warning	None	Edit the resiliency group
Datastore Cluster Deleted	Datastore cluster is deleted from the data center	6 Hours	Error	<ul style="list-style-type: none"> ■ Rehearsal ■ Migrate ■ Resync 	Edit the resiliency group
All the hosts on the applications are not reachable	All the hosts for the application are not reachable	15 minutes	Error	None	Check the connectivity with the application hosts
Application host is disconnected due to change in MAC address	Application Host is in Disconnected state	15 minutes	Error	<ul style="list-style-type: none"> ■ Rehearsal ■ Migrate 	Retry Add Host operation

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Assets does not have copy policy	Assets does not have copy policy	When vrp_host unassociated with copy policy.	Warning	None	Check if any asset has no copy policy
Backup image violates the defined RPO	Checks if the backup image violates the defined RPO	Immediate	Warning		<ul style="list-style-type: none"> ■ Check the connection state of NetBackup Notification channel. ■ Check for issues due to which backup images are not available.
CPU Usage Critical	Available compute capacity on the recovery site may be inadequate for recovering this application. This risk affects the recoverability of the services running on your production data center.	6 hours	Warning	None	Reduce the number of CPUs assigned to the virtual machines on the primary site to match the available CPU resources on failover target
Incorrect .Net version is installed	The expected .NET version is not installed or it is not compatible with the PowerShell version	2 hours	Error	<ul style="list-style-type: none"> ■ On Primary site: migrate and recover operations 	Ensure that the .NET version is installed with its compatible PowerShell version. Refer to the HSCL for compatible versions of .NET and PowerShell.
Editing the resiliency group is required	Resiliency group needs an upgrade or perform Edit operation.	Immediate	Warning	None	Edit the resiliency group using the Edit Configuration intent. Ensure that the resiliency group is online on the source datacenter before performing the edit operation

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Evacuation plan for data center has been invalidated.	Evacuation plan for data center has been invalidated, due to adding , deleting or updating a resiliency group or a VBS	Immediate	Error	None	Regenerate the evacuation plan.
Host reboot is pending after upgrade	The OS is not rebooted after upgrade operation	Immediate	Warning	None	Virtual machine requires to be rebooted after the upgrade operation
Mount point is deleted	Check if the mount point on which the assets of the resiliency group are configured, is deleted or renamed	6 hours	Error	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal 	Remount using the same mount point else you need to edit the resiliency group
PowerShell is not initialized	PowerShell is not initialized	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and rehearsal operations 	Check PowerShell Initialization on host
PowerShell is not installed	PowerShell is not installed	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and rehearsal operations 	Install PowerShell (version > 2.0) on host
Powershell Version is incorrect	Expected Powershell version not found	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and recover operations 	Install Powershell version should be 2.0 and above

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Registry Parameter LSI_SAS is not set	Registry Parameter LSI_SAS is not set	2 hours	Error	<ul style="list-style-type: none"> On Secondary site: migrate and rehearsal operations 	Change the value for registry parameter LSI_SAS->Start to 0 and refresh host discovery
Replication Gateway is not reachable	The Replication Gateway is down or not reachable from the IMS	15 minutes	Error	None	Make sure the replication gateway appliance is running and is reachable from the IMS
Virtual machine or Replication Gateway is not found in the cluster.	Virtual machine or Replication Gateway is not present in cluster using which resiliency group is created.	15 minutes	Error	migrate, rehearsal	<ul style="list-style-type: none"> If the virtual machine is moved out of the cluster or ESX server, re-add the virtual machine and perform resync operation. If the virtual machine is unregistered, perform start resiliency group operation on the resiliency group with the checkbox "Refresh storage, network, compute, and customization" as selected.
Replication state synchronizing	Data synchronization is in progress.	5 minutes	Warning	None	Wait for synchronization to complete (Replication state should be Active (Connected Consistent))

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Resync operation is pending on a resiliency group	Resync operation is pending on current resiliency group	Immediate	Error	On Secondary site: migrate operation	Execute Resync operation on current resiliency group
Resiliency group configuration drift	Disk configuration for asset(s) in the resiliency group is changed. This is a configuration drift.	2 minutes	Error	<ul style="list-style-type: none"> ■ On Primary site: rehearsal operation ■ On Secondary site: migrate, resync, and rehearsal operations 	Refresh the respective hosts, vCenter servers or Hyper-V servers and the cloud discovery. After refresh, probe the risk. If the risk still exists, remove the virtual machine from the resiliency group and re-add using the edit operation.
Resiliency group configuration error	The disk size of the virtual machine in the resiliency group has changed. This is a configuration error	2 hours	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and resync operation 	<ul style="list-style-type: none"> ■ Editing the size of a disk is not supported. Restore the disk size for a resiliency group having multiple virtual machines. ■ Edit the resiliency group by removing affected hosts and then add it again to re-protect. ■ For resiliency group having only one virtual machines delete it and recreate again.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Resiliency group outage in datacenter	Outage has been declared for the resiliency group in the datacenter	Immediate	Error	None	Perform remediation steps to clear outage in the specified datacenter. Run a Resync or Clear outage operation (as applicable) to indicate that the outage has been cleared
Data sync failed between Resiliency Manager and database.	Data sync failed between Resiliency Manager and database.	As soon as the vrp_rm vertex gets updated with property db_status as value "Data sync failed"	Error	None	Perform Resync operation for Resiliency Manager
SAN Policy Offline Shared	SAN policy on the Windows host is Offline Shared	2 hours	Warning	None	Change the SAN policy on the Windows host to Online Shared and refresh the host discovery information

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Stale configuration :: Object deleted	Asset is unavailable	As soon as discovery reports delete of addressable objects.	Error	<ul style="list-style-type: none"> ■ On Primary site: Start, stop, migrate, resync, recover, and cleanup rehearsal operations. ■ On Secondary site: Migrate, rehearsal, resync, and recover operations. 	Edit the resiliency group.
Stale configuration :: Object unreachable	Asset is unreachable	As soon as discovery reports DISCONNECTED or NOT REACHABLE fault for addressable objects.	Error	<ul style="list-style-type: none"> ■ On Primary site: Start, stop, migrate, , resync , recover, and cleanup operations. ■ On Secondary site: Migrate, rehearsal resync, and recover operations. 	Edit the resiliency group.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
The migrated virtual machine is not added to the target IMS.	The migrated virtual machine is not added to the target IMS.	45 minutes	Error	<ul style="list-style-type: none"> ■ On Secondary site: migrate and resync operation 	Refer to the documentation to know the possible reasons for failure of add host operation
Unable to get VMX	Unable to backup virtual machine configuration file	Immediate	Error	<ul style="list-style-type: none"> ■ On Primary site: rehearsal, migrate and recover operations 	Check the state of IMS, its corresponding assets such as the hypervisors and vCenter servers. Perform edit resiliency group operation.
Unable to update virtual machine configurations file	Unable to backup latest virtual machine configuration	Immediate	Error	None	Check the state of IMS, its corresponding assets such as the hypervisors and vCenter servers. Perform edit resiliency group operation.
vCenter server is removed from IMS	vCenter server is removed from IMS	Immediate	Error	<ul style="list-style-type: none"> ■ On Primary site: start, stop, rehearsal, and migrate operations ■ On Secondary site: start, stop, rehearsal, and migrate operations 	Add the vCenter server to the IMS.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
VCS Servicegroup Faulted	VCS Servicegroup is in Faulted state	1 hour	Error	None	Resolve the fault on VCS Servicegroup
Insufficient quota on target vCloud Director	Sufficient quota(CPUs/Memory/Storage) is not available on target vCloud Director.	5 minutes	Error	None	Sufficient quota should be available on the target vCloud Director
Virtual machine is deleted	One or more virtual machines are deleted or unregistered. The virtual machines belong to a resiliency group that is configured for remote recovery. This affects the recoverability of the resiliency group.	6 hours	Error	On Secondary site: migrate operation	Edit the resiliency group to remove the virtual machines that are deleted or unregistered.
Virtual machine is not protected	Virtual machine is not configured for remote recovery	Immediate	Error	None	If the virtual machine is in production data center then configure the virtual machine for remote recovery. If the virtual machine is in vCloud data center then ensure that disk.EnableUUID property is set to TRUE on the VRP_VAPP_TEMPLATE virtual machine as well as on the migrated virtual machine. After the risk is resolved, perform the Resync operation to avoid disk corruption during migrate or migrate back operation.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
VMware discovery failed	VMware discovery is failed . Unable to establish secure connection with vCenter server as the SSL/TLS handshake has failed	6 hours	Error	None	1. In case of a password change, resolve the password issue and refresh the vCenter server configuration. 2. Install valid CA certificates of vCenter server in the Resiliency Platform and refresh the vCenter server configuration.
IO Filter is not replicating the IOs from the virtual machine	IO Filter has encountered a fatal error	When IMS is receiving NOOP snmp event.	Error	<ul style="list-style-type: none"> ■ On Primary site: migrate resync deepstart (Perform start operation after reverse replication is complete) ■ On Secondary site: migrate resync deepstart (Perform start operation after reverse replication is complete) 	If IO filter has encountered errors, either invoke the edit resiliency group workflow to remove and re-add asset from the resiliency group or delete the resiliency group and create it again

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Cloud discovery failed	<p>1. Cloud discovery has failed.</p> <p>2. Unable to establish secure connection with cloud resource as the SSL/TLS handshake has failed.</p>	<p>1. After 5 minutes</p> <p>2. After 5 minutes</p>	<p>1. Error</p> <p>2. Error</p>	<ul style="list-style-type: none"> ■ On primary site: migrate, recover, rehearsal, cleanup rehearsal, and resync operations. ■ On secondary site: start, stop, migrate, and resync operations 	<p>1. Edit the cloud configuration to resolve the issue. If risk persists contact Veritas Support.</p> <p>2. Install valid CA certificate of the cloud resource in Resiliency Manager and then refresh the cloud configuration.</p>
Enclosure failed	Enclosure discovery has failed.	After 5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: migrate, recover, rehearsal, cleanup rehearsal, and resync operations. ■ On secondary site: migrate, and resync operations 	<p>General: Edit the enclosure configuration and provide valid SSH host keys. If the risk persists, contact Veritas Support.</p> <p>NetApp SSL resolution: Install valid CA certificates of NetApp enclosure in the Resiliency Platform and refresh the NetApp enclosure configuration.</p> <p>NetApp general resolution: Unable to fetch enclosure details. If the risk persists contact Veritas Support.</p>

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Cloud authentication failed	Cloud credentials are incorrect	After 5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: migrate, recover, rehearsal, cleanup rehearsal, and resync operations. ■ On secondary site: start, stop, migrate, and resync operations 	Edit cloud configuration and provide correct credentials to resolve the issue. In case of AWS, check the IAM role with proper privileges is attached to IMS.
Cloud connection timeout	Connection timed out fetching information about cloud resources.	After 5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: migrate, recover, rehearsal, cleanup rehearsal, and resync operations. ■ On secondary site: start, stop, migrate, and resync operations 	Resolve network connectivity between IMS and cloud data center and then refresh the cloud configuration.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
NTP Time Sync Failed	NTP time skew. Time skew must be less than 3 seconds.	5 minutes	Warning	None	Synchronize with NTP server.
NTP Time Unsynchronized	Not able to synchronize with the NTP server.	5 minutes	Warning	None	Synchronize with NTP server.
NTP Time Indeterminate	NTP status indeterminate	5 minutes	Warning	None	Synchronize with NTP server.
Resiliency Group Configuration Drift for Network changed of some of the assets in the Resiliency Group	This risk is raised if network of some of the assets in the Resiliency Group is changed after the Resiliency Group is created. The change can be in the VLAN, vSwitch or cloud network settings.		Error		The risk is resolved when the deleted network gets discovered in Veritas Resiliency Platform. Or the network update risk will be resolved after successful editing the Resiliency Group.
The network setting of the virtual machine 'Connect at Power On' is disabled.	The network setting of the virtual machine 'Connect at Power On' is disabled.	Immediate	Warning	None	Ensure to set the 'Connect At Power On' network adapter settings of the virtual machine is enabled.
Node added to Infoscale Cluster	A new node is added to InfoScale cluster and the node has not been configured into Resiliency Platform yet.	Raised when Resiliency Platform InfoScale cluster discovery detects the node addition in near real time after the node has been added to the cluster.	Warning	None	Perform reconfigure operation on the InfoScale cluster to configure the node in Resiliency Manager.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Appliance Storage critical	At least 85% of the disk capacity is being utilized. This risk affects the continuity of the services running on your appliance.	5 min	Warning	All	Free up the disk space or increase the disk size of the appliance.
ISO already mounted	ISO is mounted on the host	'5 min	Error	Rehearsal, Cleanup rehearsal	Remove or detach the ISO from VMware virtual machine.
Replication Gateway Service Down	One or more services on the Replication Gateway appliance are not running.	5 min	Error	<ul style="list-style-type: none"> ■ On the primary site: migrate and resync operations ■ On secondary site: migrate, recover , resync , rehearsal and, cleanup rehearsal operations. 	Make sure all the services on the Replication Gateway appliance are running.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Replication state is inactive	This risk affects the recoverability of the services running on your source data center.	10 mins	Warning	None	

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
					<p>Perform the following steps:</p> <ol style="list-style-type: none"> 1 Make sure the source Replication Gateway is powered on and all services are running. 2 In case of in-guest replication, if risk is raised: <ul style="list-style-type: none"> ■ Make sure workload is powered on.. ■ Make sure workload should communicate with source Replication Gateway over the network 3 In case of VAIO: <ul style="list-style-type: none"> ■ Make sure ESXi server on which virtual machine is residing is able to communicate with source Replication Gateway over the network. ■ In case target Replication Gateway is replaced before migrate operation and

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
					<p>after migrate operation, risk is raised on the resiliency group, perform a resync operation for resiliency group. The resync operation in this case will perform diff sync. If resync is performed on resiliency group after migrate operation, usually it performs full sync. Hence, resync needs to be done carefully for those resiliency group which has this risk after successful migrate operation.</p> <ul style="list-style-type: none"> ■ In case vMotion happens for virtual machine is being snapshotted, this risk may appear. but it gets resolved in next CG update.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
					<ul style="list-style-type: none"> ■ If a policy is detached from the workload disk, a risk appears. Performing resync operation should resolve the risk.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
					<p>4 If risk is raised on Replication Gateway after VBS migrate operation.</p> <ul style="list-style-type: none"> ■ If first attempt of VBS migrate operation fails in reverse-replication step after updating CG roles, consecutive VBS migrate operation will skip reverse replication and replication will remain inactive. Here, check source CG state on gateway. If the CG state is set to “stopped”, then perform resync operation for individual resiliency group. If resync is performed on VBS, full sync may occur for some of the resiliency group.

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
<p>Replication configuration is broken for resiliency group</p>	<p>Replication configuration is broken for resiliency group.</p>	<p>Real time after recovering the resiliency groups configured with multiple recovery point.</p>	<p>Error</p>	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal ■ Recover with replication path ■ Cleanup rehearsal ■ Start operation with network and storage refresh 	<p>Perform resync operation to repair the replication for the recovered virtual machine.</p>
<p>Protection configuration data for replication configuration is being updated for the resiliency group.</p>	<p>The protection configuration data for replication configuration is getting updated for the resiliency group.</p>	<p>Real time after recovering the resiliency groups configured with multiple recovery point.</p>	<p>Error</p>	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal ■ Recover with replication path ■ Cleanup rehearsal ■ Start operation with network and storage refresh 	<p>Once the replication configuration is updated successfully, this risk goes away.</p>

Table 2-55 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
VMWare ESXi Cluster Membership Change	Checks if any new ESX host is added to the ESX cluster.	30 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Rehearsal 	<p>To resolve the risk perform the below steps :-</p> <ol style="list-style-type: none"> 1 Create or change the required network configurations according to the environment and then Edit the Resiliency Group with Edit Configuration intent. 2 Remove the ESXi host from the cluster.
Maintenance Mode enabled on VMware ESXi host	Checks if maintenance mode is enabled on any of the ESX hosts.	30 minutes	Error / Warning	<ul style="list-style-type: none"> ■ Migrate ■ Takeover ■ Rehearsal ■ Cleanup Rehearsal 	Disable the Maintenance Mode on the ESX host or remove the ESXi host from the cluster.

See [“About risks”](#) on page 542.

About reports

Using the Veritas Resiliency Platform console, you can generate a variety of reports. The following are the broad categories under which the reports are grouped:

- **Inventory:** Reports in this category provide information about the data centers and applications, and the virtual machines that are deployed in the data centers.
- **Recovery Assessment:** This category lists the reports that are related to the disaster recovery operations such as the migrate and take over report, and the rehearsal report.

- **Risk:** This category has two reports; Current Risk and Risk History. These reports show the summary and details of all the current and historical risks that occurred in the environment.

Reports can be scoped on the data center or global. You can subscribe for a report on a daily, weekly, monthly, quarterly, or yearly basis, or on predefined days of the week, or run on demand. Reports are available in the HTML and PDF format, or as a comma-separated file (CSV) file.

You can send a report to multiple recipients by entering the email addresses separated by a comma (,) or a semicolon (;).

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 (navigation pane)

Click **Inventory**, **Recovery Assessment**, or **Risk** to expand the category.

- 2 Click **Schedule** on the desired report.
- 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 (;).

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 (navigation pane)

Click **Inventory**, **Recovery Assessment**, or **Risk** to expand the category.

2 Click **Run** on the desired report.

3 In the **Run Report** wizard panel, specify the following information, 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 (;).

Viewing reports

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

Resiliency Groups and VBS Summary	Provides details about the resiliency groups and VBSs in the data centers across all sites.
-----------------------------------	---

VM Inventory	<p>Provides the platform distribution and the OS distribution details of the virtual machines that are deployed in the data centers in the form of a pie chart.</p> <p>The Details table provides additional information for each virtual machine.</p> <p>For virtual machines on the Hyper-V Server, the report displays the total memory instead of allocated memory.</p> <p>Hyper-V virtual machines which are in offline state are displayed in the Unknown category.</p>
License Entitlement Report	<p>License Entitlement Report provides details about the licenses that are deployed in your datacenter. There are 3 types of licenses:</p> <ul style="list-style-type: none"> ■ Veritas Resiliency Platform FETB (Per-FETB) ■ Veritas Resiliency Platform Compute (Per-Core) ■ Veritas Resiliency Platform Compute (Per-VM)
Notification Throttling Report	<p>Notification Throttling Report displays all the notifications which are currently throttled and are waiting to be raised.</p>
Activity Distribution History	<p>Provides information about tasks, such as migrate, recover, rehearse, start, and stop, performed for a specified duration.</p>
Recovery Activity History by RG	<p>Provides historical information about recovery tasks, such as migrate, recover, and rehearse for each resiliency group.</p>
Recovery Activity History by VBS	<p>Provides historical information about recovery tasks, such as migrate, recover, and rehearse for each VBS.</p>
Metering	<p>Provides details of the virtualization servers that are protected for disaster recovery.</p> <p>You can view the total number of servers that are protected for disaster recovery. For these servers you can view the total memory, processor cores, and the total storage.</p>

VBS RPO

Provides Recovery Point Objective (RPO) details for all the virtual business services (VBS) in the resiliency domain.

The bar chart provides information on the top VBS with maximum RPO lag.

You can view the lag in the last replication and the replication date for all the VBS in the table.

To view a report

1 Navigate



Reports (navigation pane)

Click **Inventory**, **Recovery Assessment**, or **Risk** to expand the category.

2 Do one of the following:

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

Managing a running activity

Using the Veritas Resiliency Platform console, you can abort a task or an operation which is currently running. And you can also resume an operation which is in Pause state.

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

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

While configuring a resiliency group or a Virtual Business Service (VBS) for disaster recovery, you can select the pause or the manual intervention points. These manual intervention points let you pause the Migrate and Recover operations. You can

resume the workflow by selecting the Resume option on **Current** activities page or **Recent Activities** page.

See “[About manual intervention](#)” on page 525.

To abort an activity

1 Navigate

Do one of the following:



Activities (navigation pane). Skip to [2](#)

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

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

- Right click and select **Abort**.
- Click on the vertical ellipsis and select **Abort**
- Right click and select **Details**. Click **Abort** on the details page.

To resume a paused activity

1 Navigate

Do one of the following:



Activities (navigation pane). Skip to [2](#)

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

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

- Right click and select **Resume**.
- Click on the vertical ellipsis and select **Resume**
- Right click and select **Details**. Click **Resume** on the details page.

Miscellaneous references

After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish

menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.

- See [“About klish”](#) on page 584.
- See [“About applying updates to Resiliency Platform”](#) on page 631.
- See [“Virtual appliance security features”](#) on page 667.

About klish

Once the Veritas Resiliency Platform virtual appliance is deployed and configured, you are given limited, menu-based access to the operating system and the product. You need to use Command Line Interface Shell (klish) menu to manage the configuration-related changes to the product.

Below are the Klish options:

Table 2-56 Klish main menu options

Menu option	Description
manage	Manage the Veritas Resiliency Platform appliance
monitor	Monitor the Veritas Resiliency Platform appliance activities
network	Change some of the network configurations
settings	Change the system settings
hotfix	Manage the Veritas Resiliency Platform hotfixes
support	Access the Veritas Resiliency Platform logs
updates	Manage Veritas Resiliency Platform updates and patches
utilities	Run the miscellaneous utilities of the appliances

After the product configuration, whenever you log in to the Resiliency Platform appliance, you get the main menu of klish. This menu is the starting point, from which you can configure, manage, monitor, and support your application using the command line. You can reconfigure or modify some of the appliance settings that are configured through the product bootstrap. Following are the settings that you can reconfigure using klish:

- **Network settings:**

You can reconfigure the subnet mask, IP, default gateway, DNS server, route, traceroute, SSH-enable-NIC, NIC for accessing product user interface and search domains using the klish menu. You cannot reconfigure the hostname

that you had configured through the bootstrap process. In case of static DHCP, you cannot change the network settings using the klish menu. You cannot change the network settings for any component that is configured in the cloud environment.

- **System settings:**

You can reset the time zone, perform operations related to NTP server, shut down the appliance, reboot using the klish menu. Changing the system settings can affect the product functionality if incorrect values are set. You can also perform logical volume management (LVM) operations such as adding a disk or removing a disk using the klish menu.

- **About updates:**

You can apply patch updates on Resiliency Platform virtual appliance for update, rollback a previously prepared update and view the latest version of the Resiliency Platform. You can also configure the repository, display the current repository configuration and remove the repository.

How to use help in Klish

You can press the **tab** key to display the menu options or you can run the `help` command to get detailed help on how to use klish. Use **space** key for auto-completion of command. If you get the `Syntax Error: Illegal command line error` or `Syntax Error: The command is not completed error`, press **?** key to display detailed help on the required parameter.

Lock mode in Klish

If a klish command is expected to perform any operation on an entity such as start or stop services, it goes into lock mode and does not allow any other operation from any other session to be performed till the first operation gets completed. In such a scenario, you may encounter the following warning:

```
Warning: Try to get lock. Please wait...
```

After waiting for some time, if the operation still cannot be performed due to the lock, then you may encounter the following error:

```
Error: Can't get lock
```

In this case, you need to execute the same command after waiting for some time. The operation is performed if the lock gets released by that time.

Best practice for using Klish

At times, you may not be able to run the klish commands if the `/var/opt` directory is fully utilized and there is no space to run the klish commands. We now raise risks if disk space is running low on any appliance. Once this issue occurs, there is no way to recover from this situation. Hence, you need to periodically check if the

space in that directory is getting fully occupied, and provision for an extra disk accordingly.

See “[Klish menu options for Resiliency Manager](#)” on page 586.

See “[Klish menu options for IMS](#)” on page 598.

See “[Klish menu options for Replication Gateway](#)” on page 612.

Klish menu options for Resiliency Manager

Following are the options available for Resiliency Manager using klish menu:

Table 2-57 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
hotfix	Manage hotfixes Table 1-75
manage	Manage appliance Table 1-76
monitor	Monitor appliance activities Table 1-78
network	Manage network configuration Table 1-79
settings	Manage appliance settings Table 1-88
support	To access logs Table 1-97
updates	Manage updates and patches Table 1-101
utilities	Run miscellaneous utilities Table 1-102

Table 2-58 Options available with **hotfix** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
apply-hotfix	Apply the specified hotfix
list-applied-hotfixes	List the applied hotfixes
list-available-hotfixes	List the available hotfixes
uninstall-hotfix	Uninstall the specified hotfix

Table 2-59 Options available with **manage** command

Menu option	Description
back	Return to the previous menu
configure	Configure Resiliency Platform component or show the configured component
exit	Log out from the current CLI session
services	Manage the appliance services Table 1-77
help	Display an overview of the CLI syntax

Table 2-60 Options available with **services** command

Menu option	Description
force	Perform operations forcefully by skipping services validations. <ul style="list-style-type: none">■ force restart service name command restarts the service name mentioned forcefully.■ force stop service name command stops the service name mentioned forcefully. You can provide multiple service names (comma separated) or can provide ALL for all services

Table 2-60 Options available with **services** command (*continued*)

Menu option	Description
restart	Restart Resiliency Platform services Two options available are: <code>restart all</code> where, <i>all</i> means all the services. <code>restart service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).
start	Start Resiliency Platform services Two options available are: <code>start all</code> where, <i>all</i> means all the services. <code>start service name</code> where, <i>service name</i> 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: <code>status all</code> where, <i>all</i> means all the services. <code>status service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).
stop	Stop Resiliency Platform services Two options available are: <code>stop all</code> where, <i>all</i> means all the services. <code>stop service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).

Table 2-61 Options available with **monitor** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
top	Display the top process information

Table 2-61 Options available with **monitor** command (*continued*)

Menu option	Description
who	Display who is currently logged into the appliance
uptime	Display the uptime statistics for the appliance
FSuage	Display filesystem usage

Table 2-62 Options available with **network** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
dns	Show or set the DNS server or manage the options for resolv.conf file Table 1-80
ip	Show the IP address Table 1-82
route	View and manipulate the IP routing table Table 1-83
search-domain	Show or change the domain Table 1-84
traceroute	Trace packet routes to a particular host. You can also specify a port to trace the packet routes.
ssh-enabled-nic	Show or update SSH enabled NIC Table 1-85
nic-configuration	Show and configure the NIC Table 1-86
nic-for-UI	Show or update NICs configured to access product user interface Table 1-87

Table 2-63 Options available with **dns** command

Menu option	Description
options	Show, add, or remove options to the /etc/resolv.conf file. Refer to the documentation of resolv.conf for a list of available options and their purpose. Table 1-81
set	Configure Domain Name Server
show	Show the current Domain Name Server

Table 2-64 Options available with **options** command

Menu option	Description
add	Add a resolv.conf option
remove	Remove a resolv.conf option
show	Show options of resolv.conf file

Table 2-65 Options available with **IP** command

Menu option	Description
show	Show the current IP address

Table 2-66 Options available with **route** command

Menu option	Description
add	Set a default route or a route for a host or a subnet
delete	Delete the route entry from the routing table
show	Display your current routing table

Table 2-67 Options available with **search-domain** command

Menu option	Description
add	Add a search-domain
remove	Remove the search domain name
show	Show the search domain settings

Table 2-68 Options available with **ssh-enabled-nic** command

Menu option	Description
show	Show the NICs on which SSH is enabled. By default, SSH is enabled on all the NICs.
add	Add network interface to enable SSH on it
remove	Remove network interface to disable SSH on it

Table 2-69 Options available with **nic-configuration** command

Menu option	Description
show	Show details of NIC configuration like hostname, IPv4 or IPv6 address, prefix, gateway etc.
set	Configure the NICs which are not used while bootstrapping.

Table 2-70 Options available with **nic-for-UI** command

Menu option	Description
show	Show the NICs which are used to access product web user interface.
set	Set a NIC to access product web user interface from existing configured NICs.
remove	Remove one of the NICs used to access product web user interface.

Table 2-71 Options available with **settings** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
date	Display the current date and time for the appliance Table 1-89
lvm	Perform operations related to logical volume manager on the Appliance Table 1-90

Table 2-71 Options available with **settings** command (*continued*)

Menu option	Description
ntp	Perform operations related to NTP server
change-password	Change the admin user password for the appliance
poweroff	Shut down the appliance
reboot	Restart the appliance
timezone	Show or change the timezone for the appliance Table 1-94
password-policies	Perform operation related to password policies of administrator user for the appliance. Table 1-95

Table 2-72 Options available with **date** command

Menu option	Description
show	Show the time and date

Table 2-73 Options available with **lv** command

Menu option	Description
add-disk	Add disk to the OS or data volume. You need to attach a disk before adding it. Table 1-91
list-free-disk	List the free disks
initialize-free-disk	Initialize the newly attached free disk
list-used-disk	List the disks used by the OS or data volume Table 1-92

Table 2-73 Options available with **lvm** command (*continued*)

Menu option	Description
remove-disk	<p>Remove disk from the data volume. Remove disk operation involves migrating data from the existing disk to a new disk. You can remove a disk only after attaching a new disk with enough storage to migrate the data.</p> <p>The command first displays the list of disks being used and you need to select the disk that you want to remove. Then it displays the list of free disks where you want to migrate data and you need to select the disk. You can choose whether to initialize the new disk or not.</p> <p>It is recommended to suspend replication of all the configured Veritas Replication Sets before performing the remove disk operation.</p>
resize-logicalvolume	<p>Resize the OS or data volume for the resized data disk</p> <p>Table 1-93</p>

Note: In case you initialize the newly-added disk during add-disk or remove-disk operation, the existing data on the new disk is deleted.

Table 2-74 Options available **add-disk** command

Menu Options	Description
data-volume	Add disk to the data volume
os-volume	Add disk to the OS volume

Table 2-75 Options available with **list-used-disk** command

Menu Options	Description
data-volume	Lists disks used by the data volume
os-volume	Lists disks used by the OS volume

Table 2-76 Options available with **resize-logical-volume** command

Menu Options	Description
data-volume	Resize the data volume
os-volume	Resize the OS volume

Table 2-77 Options available with **timezone** command

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

Table 2-78 Options available with **password-policies** command

Menu option	Description
set	Modify the administrator user password policies for the appliance. Table 1-96
show	Show the administrator user password-policies for the appliance.

Table 2-79 Options available with **password-policies set** command

Menu option	Description
max-age	Modify the maximum number of days before password change is required for administrator user.
min-age	Modify the minimum number of days before password change is required for administrator user.
min-length	Modify minimum password length for the administrator user.
warning-days	Modify number of days before a warning for administrator password expiry is given.

Table 2-80 Options available with **support** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
reset-support-password	Reset the support user password to the default installation password. This option may typically be used by Veritas support.

Table 2-80 Options available with **support** command (*continued*)

Menu option	Description
shell	Open the bash shell prompt for support user
loggather	If the appliance is configured as a Resiliency Manager, then various options will be available for collecting the Resiliency Manager logs. Table 1-98

Table 2-81 Options available with **loggather** command

Menu option	Description
basic	Gather logs of Resiliency Manager excluding database and heap dumps See Table 1-99 on page 595.
full	Gather logs of Resiliency Manager with database
db	Gather database logs of Resiliency Manager
coredump	Gather heap dumps of Resiliency Manager See Table 1-100 on page 596.
cleanup coredump	Clean up all the collected loggater heap dump files of Resiliency Manager
cleanup vrp-logs	Clean up all the collected loggater log files of Resiliency Manager
cleanup all	Clean up all the collected loggater files (vrp-logs and coredump) of Resiliency Manager
show	Lists all the loggater URLs ordered by date and time of Resiliency Manager

Table 2-82 Options available with **basic** command

Menu option	Description
Number of days	Displays the basic logs from the days (1-99) mentioned

Table 2-83 Options available with **coredump** command

Menu option	Description
days	Displays the coredump logs from the days (1-99) mentioned

Table 2-84 Options available with **updates** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
Show-version	Show the appliance version
prepare-for-update	Save the virtual appliance configuration in preparation for upgrade
rollback-update	To rollback the prepare for update operation

Table 2-85 Options available with **utilities** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
clear	Clear the screen
unmount-cd-rom	Unmount the CD-ROM from the appliance
troubleshoot run-tool	Use the troubleshooting menu options Table 1-103
vmware-tools	Perform VMware Tools operations (install, uninstall, and show-version) Table 1-104
sftp-session	Use SFTP session for file transfer operation on the SFTP server Table 1-105

Table 2-85 Options available with **utilities** command (*continued*)

Menu option	Description
azure-waagent-service	Perform Azure waagent service operation. Applicable only in Azure environment Table 1-109

Table 2-86 Options available with **troubleshoot run-tool** command

Menu option	Description
view-logs	View log files on any virtual appliance
check-port	Verify required open ports on Veritas Resiliency Platform VSA for communication with other appliance using admin password.

Table 2-87 Options available with **vmware-tools** command

Menu option	Description
install	Install the VMware Tools mounted on CD-ROM of the appliance
show-version	Show the installed version of VMware Tools on the appliance
uninstall	Uninstall the VMware Tools from the appliance

Table 2-88 Options available with **sftp-session** command

Menu option	Description
start	To start the SFTP server session Table 1-106
show-details	View the current SFTP user and session details
stop	To stop the SFTP server session

Table 2-89 Options available with **start** command

Menu option	Description
get	View the file types that can be downloaded from the SFTP server Table 1-107

Table 2-89 Options available with start command (*continued*)

Menu option	Description
put	View the file types that can be uploaded on the SFTP server Table 1-108

Table 2-90 Options available with **get** command

Menu option	Description
logs	Download the log files and directories from the SFTP server
heap-dump	Download the heap dump files of the service available on Resiliency Manager

Table 2-91 Options available with **put** command

Menu option	Description
patch	Upload the private patch on the SFTP server

Table 2-92 Options available with **azure-waagent-service** command

Menu option	Description
start	Start Azure waagent service
stop	Stop Azure waagent service
status	Show current status of Azure waagent service

Klish menu options for IMS

Following are the options available for IMS using klish menu:

Table 2-93 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

Table 2-93 Options available in the **main** menu (*continued*)

Menu option	Description
hotfix	Manage hotfixes Table 1-111
manage	Manage appliance Table 1-112
monitor	Monitor appliance activities Table 1-116
network	Manage network configuration Table 1-117
settings	Manage appliance settings Table 1-126
support	To access logs Table 1-135
updates	Manage updates and patches Table 1-138
utilities	Run miscellaneous utilities Table 1-140

Table 2-94 Options available with **hotfix** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
apply-hotfix	Apply the specified hotfix
list-applied-hotfixes	List the applied hotfixes
list-available-hotfixes	List the available hotfixes
uninstall-hotfix	Uninstall the specified hotfix

Table 2-95 Options available with **manage** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
configure	Configure Resiliency Platform component or show the configured component Table 1-113
infra-appliance	List or remove Replication Gateway appliance Table 1-114
services	Manage the appliance services Table 1-115
show-resiliency-domain	Show details of resiliency domain to which the Infrastructure Management Server is configured. The option also lists the configured Resiliency Manager(s) in the domain.

Table 2-96 Options available with **configure** command

Menu option	Description
ims_register	Register the IMS using the registration URL obtained after initiating the Add IMS operation This option is available only for an IMS appliance Add link to Add an IMS topic
show	Show the configured component

Table 2-97 Options available with **infra-appliance** command

Menu option	Description
list	List Resiliency Platform infrastructure appliance.
remove	Remove the Replication Gateway appliance. You need to remove the Gateway pair before you remove the Gateway. Remove replication gateway link topic

Table 2-98 Options available with **services** command

Menu option	Description
restart	Restart Resiliency Platform services Two options available are: <code>restart all</code> where, <i>all</i> means all the services. <code>restart service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).
start	Start Resiliency Platform services Two options available are: <code>start all</code> where, <i>all</i> means all the services. <code>start service name</code> where, <i>service name</i> 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: <code>status all</code> where, <i>all</i> means all the services. <code>status service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).
stop	Stop Resiliency Platform services Two options available are: <code>stop all</code> where, <i>all</i> means all the services. <code>stop service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).

Table 2-99 Options available with **monitor** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
top	Display the top process information

Table 2-99 Options available with **monitor** command (*continued*)

Menu option	Description
who	Display who is currently logged into the appliance
uptime	Display the uptime statistics for the appliance
check-cim-status	Display the status of Common Information Model (CIM)
FSuage	Display filesystem usage

Table 2-100 Options available with **network** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
dns	Show or set the DNS server or manage the options for resolv.conf file Table 1-118
ip	Show the IP address Table 1-120
route	View and manipulate the IP routing table Table 1-121
search-domain	Show or change the domain Table 1-122
traceroute	Trace packet routes to a particular host. You can also specify a port to trace the packet routes.
ssh-enabled-nic	Show or update SSH enabled NIC Table 1-123
nic-configuration	Show and configure the NIC Table 1-124
nic-for-UI	Show or update NICs configured to access product user interface Table 1-125

Table 2-101 Options available with **dns** command

Menu option	Description
options	Show, add, or remove options to the <code>/etc/resolv.conf</code> file. Refer to the documentation of <code>resolv.conf</code> for a list of available options and their purpose. Table 1-119
set	Configure Domain Name Server
show	Show the current Domain Name Server

Table 2-102 Options available with **options** command

Menu option	Description
add	Add a <code>resolv.conf</code> option
remove	Remove a <code>resolv.conf</code> option
show	Show options of <code>resolv.conf</code> file

Table 2-103 Options available with **IP** command

Menu option	Description
show	Show the current IP address

Table 2-104 Options available with **route** command

Menu option	Description
add	Set a default route or a route for a host or a subnet
delete	Delete the route entry from the routing table
show	Display your current routing table

Table 2-105 Options available with **search-domain** command

Menu option	Description
add	Add a search-domain
remove	Remove the search domain name
show	Show the search domain settings

Table 2-106 Options available with **ssh-enabled-nic** command

Menu option	Description
show	Show the NICs on which SSH is enabled. By default, SSH is enabled on all the NICs
add	Add NIC to enable SSH on it
remove	Remove NIC to disable SSH on it

Table 2-107 Options available with **nic-configuration** command

Menu option	Description
show	Show details of NIC configuration like hostname, IPv4 or IPv6 address, prefix, gateway etc.
set	Configure the NICs which are not used while bootstrapping.

Table 2-108 Options available with **nic-for-UI** command

Menu option	Description
show	Show the NICs which are used to access product web user interface.
set	Set a NIC to access product web user interface from existing configured NICs.
remove	Remove one of the NICs used to access product web user interface.

Table 2-109 Options available with **settings** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
date	Display the current date and time for the appliance Table 1-127
lvm	Perform operations related to logical volume manager on the Appliance Table 1-128

Table 2-109 Options available with **settings** command (*continued*)

Menu option	Description
ntp	Perform operations related to NTP server
change-password	Change the admin user password for the appliance
poweroff	Shut down the appliance
reboot	Restart the appliance
timezone	Show or change the timezone for the appliance Table 1-132
password-policies	Perform operation related to password policies of administrator user for the appliance. Table 1-133

Table 2-110 Options available with **date** command

Menu option	Description
show	Show the time and date

Table 2-111 Options available with **lv** command

Menu option	Description
add-disk	Add disk to the OS or data volume. You need to attach a disk before adding it. Table 1-129
list-free-disk	List the free disks
initialize-free-disk	Initialize the newly attached free disk
list-used-disk	List the disks used by the OS or data volume. Table 1-130

Table 2-111 Options available with **lvm** command (*continued*)

Menu option	Description
remove-disk	<p>Remove disk from the data volume. Remove disk operation involves migrating data from the existing disk to a new disk. You can remove a disk only after attaching a new disk with enough storage to migrate the data.</p> <p>The command first displays the list of disks being used and you need to select the disk that you want to remove. Then it displays the list of free disks where you want to migrate data and you need to select the disk. You can choose whether to initialize the new disk or not.</p> <p>It is recommended to suspend replication of all the configured Veritas Replication Sets before performing the remove disk operation.</p>
resize-logicalvolume	<p>Resize the OS or data volume for resized data disk</p> <p>Table 1-131</p>

Note: In case you initialize the newly-added disk during add-disk or remove-disk operation, the existing data on the new disk is deleted.

Table 2-112 Options available with **add-disk** command

Menu options	Description
data-volume	Add disk to the data volume
os-volume	Add disk to the OS volume

Table 2-113 Options available with **list-used-disk** command

Menu options	Description
data-volume	Lists disks used by the data volume
os-volume	Lists disks used by the OS volume

Table 2-114 Options with available with **resize-logicalvolume** command

Menu options	Description
data-volume	Resize the data volume
os-volume	Resize the OS volume

Table 2-115 Options available with **timezone** command

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

Table 2-116 Options available with **password-policies** command

Menu option	Description
set	Modify the administrator user password policies for the appliance.
show	Show the administrator user password-policies for the appliance.

Table 2-117 Options available with **password-policies set** command

Menu option	Description
max-age	Modify the maximum number of days before password change is required for administrator user.
min-age	Modify the minimum number of days before password change is required for administrator user.
min-length	Modify minimum password length for the administrator user.
warning-days	Modify number of days before a warning for administrator password expiry is given.

Table 2-118 Options available with **support** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
reset-support-password	Reset the support user password to the default installation password. This option may typically be used by Veritas support.
shell	Open the bash shell prompt for support user

Table 2-118 Options available with **support** command (*continued*)

Menu option	Description
loggather	If the appliance is configured as an IMS, then various options will be available for collecting the IMS logs. Table 1-136

Table 2-119 Options available with **loggather** command

Menu option	Description
basic	Gather logs of IMS without database See Table 1-137 on page 608.
full	Gather logs of IMS with database
cleanup	Clean up the loggater files of IMS
show	Lists all the loggater URLs ordered by date and time of IMS

Table 2-120 Options available with **basic** command

Menu option	Description
Number of days	Displays the basic logs from the days (1-99) mentioned

Table 2-121 Options available with **updates** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
Show-version	Show the appliance version
extra-drivers	Shows the list of the drivers and updates the extra drivers. Table 1-139
prepare-for-upgrade	Save the virtual appliance configuration in preparation for upgrade
rollback-update	Roll back the prepare for update operation

Table 2-122 Parameters needed for **extra-drivers** command

Menu option	Description
list	Show list of drivers and whether the updates available for that driver.
update	Update the drivers

Table 2-123 Options available with **utilities** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
clear	Clear the screen
unmount-cd-rom	Unmount the CD-ROM from the appliance
troubleshoot run-tool	Use the troubleshooting menu options Table 1-141
vmware-tools	Perform VMware Tools operations (install, uninstall, and show-version) Table 1-145
sftp-session	To start the SFTP server session Table 1-146
azure-waagent-service	Perform Azure waagent service operation. Applicable only in Azure environment Table 1-150
svc-delete-vdisk-timeout	To set or unset IBM SVC delete vdisk timeout value, which is the maximum amount of time that will be spent in retrying the delete vdisk operation during Cleanup Rehearsal. Table 1-151

Table 2-124 Options available with **troubleshoot run-tool** command

Menu option	Description
manage-nbu-primary-server-certificates	Manage the NetBackup certificates in Resiliency Platform Table 1-142
view-logs	View log files on any virtual appliance
check-port	Verify required open ports on Veritas Resiliency Platform VSA for communication with other appliance using admin password.

Table 2-125 Options available with **manage-nbu-primary-server-certificates** command

Menu option	Description
add	Add the certificate to the specified NetBackup primary server with token. Table 1-143
delete	Delete the certificate with specified fingerprint. Table 1-144
show	Show all the certificates.
help	Display the help text.

Table 2-126 Options available with **add** command

Menu option	Description
primary	Provide the NetBackup primary server hostname for reissuing the certificate
token	Provide the token value for the certificate registration

Table 2-127 Options available with **delete** command

Menu option	Description
fingerprint	Provide the fingerprint of specific certificate as an input to delete the certificate

Table 2-128 Options available with **vmware-tools** command

Menu option	Description
install	Install the VMware Tools mounted on CD-ROM of the appliance
show-version	Show the installed version of VMware Tools on the appliance
uninstall	Uninstall the VMware Tools from the appliance

Table 2-129 Options available with **sftp-session** command

Menu option	Description
start	To start the SFTP server session Table 1-147
show-details	View the current SFTP user and session details
stop	To stop the SFTP server session

Table 2-130 Options available with **start** command

Menu option	Description
get	View the file types that can be downloaded from the SFTP server Table 1-148
put	View the file types that can be uploaded on the SFTP server Table 1-149

Table 2-131 Options available with **get** command

Menu option	Description
logs	View the log files and directories from the SFTP server
heap-dump	Download the heap dump files of the service available on Resiliency Manager

Table 2-132 Options available with **put** command

Menu option	Description
patch	Upload the private patch on the SFTP server

Table 2-133 Options available with **azure-waagent-service** command

Menu option	Description
start	Start Azure waagent service
stop	Stop Azure waagent service
status	Show current status of Azure waagent service

Table 2-134 Options available with **svc-delete-vdisk-timeout** command

Menu option	Description
set	Sets SVC delete vdisk timeout value.
show	Shows SVC delete vdisk timeout value.
unset	Unsets SVC delete vdisk timeout value.

Klish menu options for Replication Gateway

Following are the options available for Replication Gateway using klish menu:

Table 2-135 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
hotfix	Manage hotfixes Table 1-153
manage	Manage appliance Table 1-154
monitor	Monitor appliance activities Table 1-160

Table 2-135 Options available in the **main** menu (*continued*)

Menu option	Description
network	Manage network configuration Table 1-163
settings	Manage appliance settings Table 1-172
support	To access logs Table 1-181
utilities	Run miscellaneous utilities Table 1-184

Table 2-136 Options available with **hotfix** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
apply-hotfix	Apply the specified hotfix
list-applied-hotfixes	List the applied hotfixes
list-available-hotfixes	List the available hotfixes
uninstall-hotfix	Uninstall the specified hotfix

Table 2-137 Options available with **manage** command

Menu option	Description
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
datamover	Manage Resiliency Platform Data Mover activities and objects Table 1-155
services	Manage the appliance services Table 1-157

Table 2-137 Options available with **manage** command (*continued*)

Menu option	Description
staging-storage	Perform Staging Storage Operations Table 1-158
cdp-storage	Manage CDP related storage Table 1-159

Table 2-138 Options available with **datamover** command

Menu option	Description
start	Start a Veritas Replication Set
abort	Stop a Veritas Replication Set
delete	Delete a Veritas Replication Set
resume	Resume a Veritas Replication Set
pause	Pause a Veritas Replication Set locally.
suspend-force	Suspend a Veritas Replication Set
clear-admin-wait	Clear the admin Wait status for the Veritas Replication Set
modify-quota-size	<p>Modify the size of the quota of a Veritas Replication Set. Modifying the quota affects the number of hosts that are protected with the gateway.</p> <p>You need to configure same quota size on all the peer gateways.</p> <p>Default quota size is 8000 MB. The minimum allowed quota size in direct mode is 2000 MB and in ObjectStore mode is 3500 MB. The maximum allowed quota size in direct and ObjectStore mode is 8000 MB.</p>
modify-updateset-size	<p>Modify the size of the UpdateSet. You need to configure same UpdateSet size on all the peer gateways.</p> <p>Default UpdateSet size is 500 MB. The minimum allowed UpdateSet size is 500 MB and the maximum allowed UpdateSet size is 2000 MB.</p>

Table 2-138 Options available with **datamover** command (*continued*)

Menu option	Description
modify-replication-frequency	<p>Modify the replication frequency.</p> <p>The default frequency is 120 sec.</p> <p>The minimum replication frequency allowed is 60 sec and maximum replication frequency allowed in 300 sec</p>
modify-tcpadvwin	<p>Modify the TCP socket's send and receive size between workload and the Replication Gateway. The default tcpadvwin value is 2 MB and maximum is 4MB. After the sizes are changed, replication of the consistency group has to be suspended and resumed to reflect this change.</p> <p>To suspend and resume the changes, use <code>suspend-force</code> and <code>resume</code> commands.</p> <p>Note: This command is not applicable for in-guest Windows workloads.</p>

Table 2-139 Options available with **fips** command

Menu option	Description
enable	Disable FIPS on the Replication Gateway appliance.
disable	Enable FIPS on the Replication Gateway appliance.
status	Show the current status of FIPs mode for the Replication Gateway appliance.

Table 2-140 Options available with **services** command

Menu option	Description
restart	<p>Restart Resiliency Platform services</p> <p>Two options available are:</p> <p><code>restart all</code> where, <i>all</i> means all the services.</p> <p><code>restart service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).</p>

Table 2-140 Options available with **services** command (*continued*)

Menu option	Description
start	<p>Start Resiliency Platform services</p> <p>Two options available are:</p> <p><code>start all</code> where, <i>all</i> means all the services.</p> <p><code>start service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).</p>
status	<p>Check the status of Resiliency Platform services</p> <p>Two options available are:</p> <p><code>status all</code> where, <i>all</i> means all the services.</p> <p><code>status service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).</p>
stop	<p>Stop Resiliency Platform services</p> <p>Two options available are:</p> <p><code>stop all</code> where, <i>all</i> means all the services.</p> <p><code>stop service name</code> where, <i>service name</i> is the name of a particular service. You can provide multiple service names (comma separated).</p>

Table 2-141 Options available with **staging-storage** command

Menu option	Description
add-disk	Add disk to the staging-storage. You need to attach a disk before adding it.
list-used-disk	List the disks used in staging-storage.

Table 2-141 Options available with **staging-storage** command (*continued*)

Menu option	Description
remove-disk	<p>Remove disk from the staging-storage. Remove disk operation involves migrating data from the existing disk to a new disk. You can remove a disk only after attaching a new disk with enough storage to migrate the data.</p> <p>The command first displays the list of disks being used and you need to select the disk that you want to remove. Then it displays the list of free disks where you want to migrate data and you need to select the disk. You can choose whether to initialize the new disk or not.</p> <p>It is recommended to suspend replication of all the configured Veritas Replication Sets before performing the remove disk operation.</p>
resize-logicalvolume	Resize the volume used in staging-storage.

Table 2-142 Options available with **cdp-storage** command

Menu option	Description
add-disk	To add a disk to the CDP storage
create-cdp-storage	To create a CDP storage
list-cdp-storage	To list all the CDP storage
list-used-disk	To list the disk used in the CDP storage
recover-cdp-storage	To recover the CDP storage
remove-cdp-storage	To remove the CDP storage
resize-logicalvolume	To resize the volume used in the CDP storage
recover-cdp-storage	To recover the CDP storage. This is also trigger data integrity test for the recovered CDP storage.
data-integrity-test	Perform data integrity test on data residing on respective CDP storage associated with Replication Set. Ensure that replication is not active for the opted replication set. One can perform Pause Table 1-155 before issuing this test.

Table 2-143 Options available with **monitor** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
top	Display the top process information
who	Display who is currently logged into the appliance
uptime	Display the uptime statistics for the appliance
FSuage	Display filesystem usage
datamover	Display VRP Datamover activities and objects Table 1-161

Table 2-144 Options available with **datamover** command

Menu option	Description
cdp-recovery-points	Display Veritas Replication Set details like recovery point ID, state, last received input /outputs, and time stamps
cdp-usage-stats	Displays details for each Veritas Replication Set like allocated space, free space, usage, etc.
usage-stats	Display the Resiliency Platform Data Mover usage stats
repl-sets	Display the details about Veritas Replication Sets including RPO, connection state, replication state. Display time required to sync the data and percentage of synced or replicated data. Table 1-161
update-sets	Display the list of current update sets which are in transit Table 1-161
ingress-data	Display the IO statistics for the data transfer from protected virtual or physical machine to Gateway (IOReceiver statistics) Table 1-161

Table 2-144 Options available with **datamover** command (*continued*)

Menu option	Description
network-data	<p>Display the network related statistics for data transfer between production site Gateway and recovery site Gateway (Transceiver statistics)</p> <p>Table 1-161</p>
disk-data	<p>Display the IO statistics for the data write on recovery site disks (Applier statistics)</p> <p>Table 1-161</p>
pair-status	<p>Displays information and connectivity status of the local and the peer Gateway</p> <p>Table 1-161</p>

Table 2-145 More information on the command

Datamover command	More information
repl-sets	<p>Name: Veritas Replication Set name for the protected virtual machine</p> <p>VRS-ID: Veritas Replication Set unique ID of a protected virtual machine</p> <p>Role: Role of the data center for the current Veritas Replication Set</p> <p>Data State: Replication data state for the current VRS-ID</p> <p>State: Replication State for the current VRS-ID</p> <p>Host Connection: Connection state of the protected virtual machine with the Replication Gateway</p> <p>Disks: Number of replicating disks for the protected virtual machine</p> <p>Lag (seconds): The time difference in seconds between a write operation occurs on the protected host on source side and the same gets replicated on the target side</p> <p>Admin Intervention: A flag notation if replication is broken. Check Admin wait state code</p> <p>Peer Gateway IP: IP address of the paired gateway</p>

Table 2-145 More information on the command (*continued*)

Datamover command	More information
update-sets	<p>Name: Veritas Replication Set name for the protected virtual machine</p> <p>VRS-ID: Veritas Replication Set unique ID of a protected virtual machine</p> <p>USID: Unique ID of the current update set. This is an increasing counter for each update set</p> <p>State: Current state of the Replication Update Set</p> <p>Size: Size of the data which is replicated in this update set</p> <p>Elapsed Time: Time of Update Set in the current state</p>
ingress-data	<p>VRS-ID: Veritas Replication Set Unique ID of a protected virtual machine</p> <p>USID: Unique ID of the current update set</p> <p>#Disk: Number of replicating disk or disks for protected virtual machine</p> <p>State: Shows different states of the virtual machine which is attached to the Replication Gateway. (INIT/DISCONNECTED/ACTIVE/SUSPENDED/ABORTED)</p> <p>Rate: IO rate of data from the protected host to Replication Gateway (Example: For the current update set, 907.0MB data has been received at the rate of 403.1Mb/s)</p> <p>Latency: This depicts the latency between the protected host to source Replication Gateway. Latency gives information about which component is slower</p> <p>Local deduplication: To know how much data is sent and the old data is cancelled (Amount of data written will be cancelled / total data size)</p>

Table 2-145 More information on the command (*continued*)

Datamover command	More information
network-data	<p>VRS-ID: Veritas Replication Set Unique ID of a protected virtual machine</p> <p>USID: Unique ID of the current update set</p> <p>Direction:</p> <p>Send: If update set is being sent from the Replication Gateway.</p> <p>Receive: If update set is being received by the Replication Gateway.</p> <p>Size: Size of the replicated data in this update set</p> <p>Rate: This is the rate at which the data is being written to the target disks by replication gateway</p> <p>Latency: Latency gives information about which component is slower</p> <p>Compression Ratio: Shows how much data is compressed while sending over WAN. The compression ratio is equal to actual data in update set/data sent over WAN</p>
disk-data	<p>VRS-ID: Veritas Replication Set Unique ID of a protected virtual machine</p> <p>USID: Unique ID of the current update set</p> <p>#Disk: Number of replicating disk or disks for protected virtual machine</p> <p>Size: Size of the replicated data in this update set</p> <p>Rate: This is the rate at which the data is being written to the target disks by replication gateway</p> <p>(Example: If update is 907MB, then rate at which it is written to the disk can be 403.1Mbps)</p> <p>Latency (Read, Write): Shows relative latency between staging area disk read and target disk write</p>

Table 2-145 More information on the command (*continued*)

Datamover command	More information
pair-status	<p>Pair Name: Name of the pair created between local and peer gateway</p> <p>Peer Gateway IP: IP address of the paired gateway</p> <p>Gateway Pair ID: Unique ID of the gateway which is paired</p> <p>Peer Gateway ID: Unique ID of the peer gateway</p> <p>Config State: Configuration state of the pair</p> <p>Connection State: Status of the connection between the local and the peer gateway</p>

Table 2-146 Options available with **network** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
dns	<p>Show or set the DNS server or manage the options for resolv.conf file</p> <p>Table 1-164</p>
ip	<p>Show the IP address</p> <p>Table 1-166</p>
route	<p>View and manipulate the IP routing table</p> <p>Table 1-167</p>
search-domain	<p>Show or change the domain</p> <p>Table 1-168</p>
traceroute	Trace packet routes to a particular host. You can also specify a port to trace the packet routes.
ssh-enabled-nic	<p>Show or update SSH enabled NIC</p> <p>Table 1-169</p>
nic-configuration	<p>Show and configure the NIC</p> <p>Table 1-170</p>

Table 2-146 Options available with **network** command (*continued*)

Menu option	Description
nic-for-UI	Show or update NICs configured to access product user interface Table 1-171

Table 2-147 Options available with **dns** command

Menu option	Description
options	Show, add, or remove options to the <code>/etc/resolv.conf</code> file. Refer to the documentation of <code>resolv.conf</code> for a list of available options and their purpose. Table 1-165
set	Configure Domain Name Server
show	Show the current Domain Name Server

Table 2-148 Options available with **options** command

Menu option	Description
add	Add a <code>resolv.conf</code> option
remove	Remove a <code>resolv.conf</code> option
show	Show options of <code>resolv.conf</code> file

Table 2-149 Options available with **IP** command

Menu option	Description
show	Show the current IP address

Table 2-150 Options available with **route** command

Menu option	Description
add	Set a default route or a route for a host or a subnet
delete	Delete the route entry from the routing table
show	Display your current routing table

Table 2-151 Options available with **search-domain** command

Menu option	Description
add	Add a search-domain
remove	Remove the search domain name
show	Show the search domain settings

Table 2-152 Options available with **ssh-enabled-nic** command

Menu option	Description
show	Show the NICs on which SSH is enabled. By default, SSH is enabled on all the NICs.
add	Add NIC to enable SSH on it
remove	Remove NIC to disable SSH on it

Table 2-153 Options available with **nic-configuration** command

Menu option	Description
show	Show details of NIC configuration like hostname, IPv4 or IPv6 address, prefix, gateway etc.
set	Configure the NICs which are not used while bootstrapping.

Table 2-154 Options available with **nic-for-UI** command

Menu option	Description
show	Show the NICs which are used to access product web user interface.
set	Set a NIC to access product web user interface from existing configured NICs.
remove	Remove one of the NICs used to access product web user interface.

Table 2-155 Options available with **settings** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session

Table 2-155 Options available with **settings** command (*continued*)

Menu option	Description
help	Display an overview of the CLI syntax
date	Display the current date and time for the appliance Table 1-173
lvm	Perform operations related to logical volume manager on the appliance Table 1-174
ntp	Perform operations related to NTP server
change-password	Change the admin user password for the appliance
poweroff	Shut down the appliance
fips	Enable, disable, or view the status of FIPS mode for a Replication Gateway Table 1-156
configure	Configure Resiliency Platform component or show the configured component
reboot	Restart the appliance
timezone	Show or change the timezone for the appliance Table 1-178
password-policies	Perform operation related to password policies of administrator user for the appliance Table 1-179

Table 2-156 Options available with **date** command

Menu option	Description
show	Show the time and date

Table 2-157 Options available with **lvm** command

Menu option	Description
add-disk	Add disk to the OS or data volume. You need to attach a disk before adding it. Table 1-175
list-free-disk	List the free disks
initialize-free-disk	Initialize the newly attached free disk
list-used-disk	List the disks used by the OS or data volume. Table 1-176
remove-disk	Remove disk from the data volume. Remove disk operation involves migrating data from the existing disk to a new disk. You can remove a disk only after attaching a new disk with enough storage to migrate the data. The command first displays the list of disks being used and you need to select the disk that you want to remove. Then it displays the list of free disks where you want to migrate data and you need to select the disk. You can choose whether to initialize the new disk or not. It is recommended to suspend replication of all the configured Veritas Replication Sets before performing the remove disk operation.
resize-logicalvolume	Resize the OS or data volume for resized data disk. Table 1-177

Table 2-158 Options available with **add-disk** command

Menu options	Description
data-volume	Add disk to the data volume
os-volume	Add disk to the OS volume

Table 2-159 Options available with **list-used-disk** command

Menu options	Description
data-volume	Lists disks used by the data volume
os-volume	Lists disks used by the OS volume

Table 2-160 Options available with **resize-logical-volume** command

Menu options	Description
data-volume	Resize the data volume
os-volume	Resize the OS volume

Note: In case you initialize the newly-added disk during add-disk or remove-disk operation, the existing data on the new disk is deleted.

Table 2-161 Options available with **timezone** command

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

Table 2-162 Options available with **password-policies** command

Menu option	Description
set	Modify the administrator user password policies for the appliance.
show	Show the administrator user password-policies for the appliance.

Table 2-163 Options available with **password-policies set** command

Menu option	Description
max-age	Modify the maximum number of days before password change is required for administrator user.
min-age	Modify the minimum number of days before password change is required for administrator user.
min-length	Modify minimum password length for the administrator user.
warning-days	Modify number of days before a warning for administrator password expiry is given.

Table 2-164 Options available with **support** command

Menu option	Description
back	Return to the previous menu
exit	Log out from the current CLI session
help	Display an overview of the CLI syntax
reset-support-password	Reset the support user password to the default installation password. This option may typically be used by Veritas support.
shell	Open the bash shell prompt for support user
loggather	If the appliance has been configured as a Replication Gateway, then loggather full command will collect the logs of the Replication Gateway. Table 1-182

Table 2-165 Options available with **loggather** command

Menu option	Description
full	Gather logs of Replication Gateway with databaseSee Table 1-183 on page 628.
cleanup	Clean up the loggather files of Replication Gateway
show	Lists all the loggather URLs ordered by date and time of Replication Gateway

Table 2-166 Options available with **full** command

Menu option	Description
Number of days	Displays the full logs from the days (1-99) mentioned

Table 2-167 Options available with **utilities** command

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

Table 2-167 Options available with **utilities** command (*continued*)

Menu option	Description
clear	Clear the screen
unmount-cd-rom	Unmount the CD-ROM from the appliance
troubleshoot run-tool	Use the troubleshoot menu options Table 1-185
vmware-tools	Perform VMware Tools operations (install, uninstall, and show-version) Table 1-186
sftp-session	Use SFTP session for file transfer operation on the SFTP server Table 1-187
azure-waagent-service	Perform Azure waagent service operation. Applicable only in Azure environment
device-path-id	Show details of the attached disks to the Replication Gateway Table 1-188

Table 2-168 Options available with **troubleshoot run-tool** command

Menu option	Description
view-logs	View log files on any virtual appliance
check-port	Verify required open ports on Veritas Resiliency Platform VSA for communication with other appliance using admin password.

Table 2-169 Options available with **vmware-tools** command

Menu option	Description
install	Install the VMware Tools mounted on CD-ROM of the appliance
show-version	Show the installed version of VMware Tools on the appliance
uninstall	Uninstall the VMware Tools from the appliance

Table 2-170 Options available with **sftp-session** command

Menu option	Description
start	To start the SFTP server session Table 1-189
show-details	View the current SFTP user and session details
stop	To stop the SFTP server session

Table 2-171 Options available with **device-path-id** command

Menu option	Description
size	Show device path, disk ID, and size of the attached disks to the Replication Gateway

Table 2-172 Options available with start command

Menu option	Description
get	View the file types that can be downloaded from the SFTP server Table 1-190
put	View the file types that can be uploaded on the SFTP server Table 1-191

Table 2-173 Options available with **get** command

Menu option	Description
logs	Download the log files and directories from the SFTP server
heap-dump	Download the heap dump files of the service available on Resiliency Manager

Table 2-174 Options available with **put** command

Menu option	Description
patch	Upload the private patch on the SFTP server

Table 2-175 Options available with **azure-waagent-service** command

Menu option	Description
start	Start Azure waagent service
stop	Stop Azure waagent service
status	Show current status of Azure waagent service

About applying updates to Resiliency Platform

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

Updating the Resiliency Platform (Resiliency Manager and IMS) involves saving configuration to the data disk of the previous version appliance and then attaching the data disk to a new, freshly deployed virtual appliance.

For more details, refer the below topics:

About deploying the Resiliency Platform virtual appliances

Deployment workflows See [“Deployment workflows”](#) on page 370.

In Veritas Resiliency Platform, you can apply updates to the following:

- Veritas Resiliency Platform virtual appliance
- Veritas Resiliency Platform add-ons
- Host packages on the assets that are added to the Infrastructure Management Server (IMS) as a host
- Veritas Replication VIB

Upgrade path for Veritas Resiliency Platform 10.0

Veritas Resiliency Platform 3.6 is the minimum version supported for upgrade to version 10.0. If your installed version is older than minimum supported version, then you should first upgrade to your supported upgrade version.

Appliance upgrade sequence for Veritas Resiliency Platform 10.0

Appliances should be upgraded in following sequence:

1. All Resiliency Manager appliances
2. All IMS appliances
3. All Replication Gateway appliances

4. All managed hosts

Table 2-176 Upgrade path for Veritas Resiliency Platform 10.0

Current version	Can upgrade to
v3.6.0.0 with any hotfix applied	v10.0
v4.0.0.0 with any hotfix applied	v10.0

Considerations for applying update

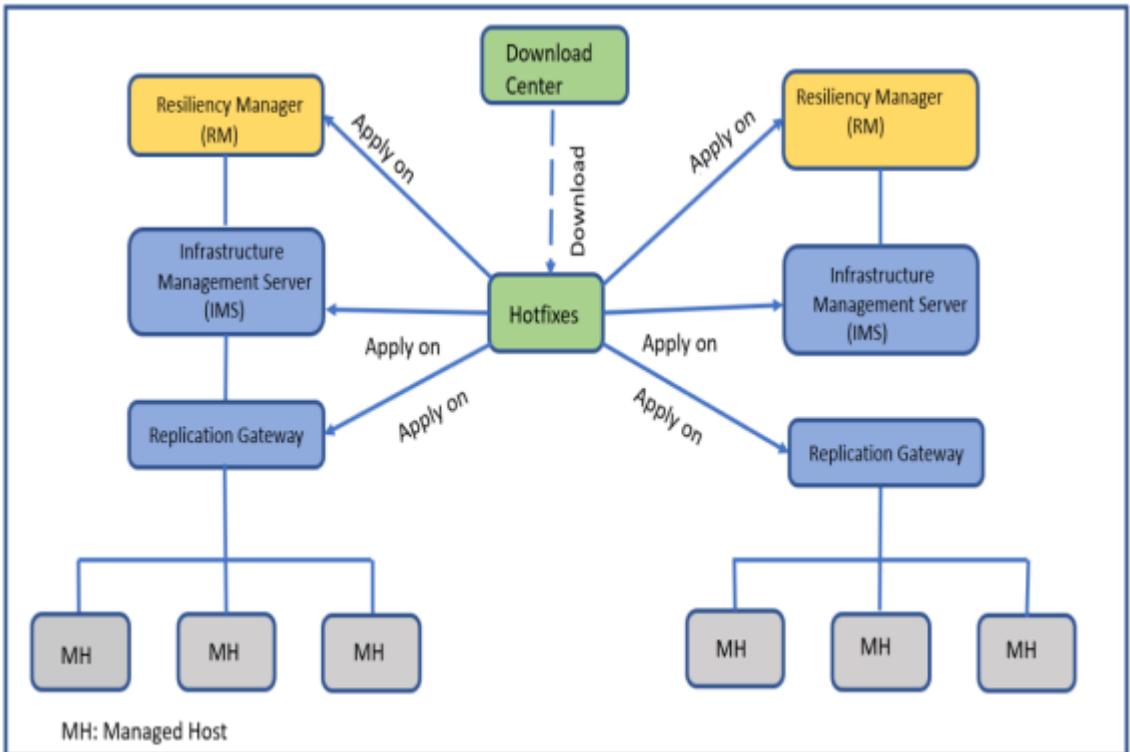
Following are some considerations for applying update to the Resiliency Platform components:

- The minimum version of the virtual appliances should be 3.6 or 4.0 to upgrade to Resiliency Platform 10.0.
- You must apply the updates on Resiliency Manager and IMS to take the complete advantage of the changes available in the updates. You need to replace the Replication Gateway to reflect the new upgrade changes in version 4.0.
- Ensure that no workflow or resiliency plans are running in the background.
- After applying updates, you can access the Resiliency Manager web user interface through the NIC that is used for communication with other Resiliency Manager.

You can change the NICs that are used to access product user interface. Use `network > nic-for-UI` Klish command to add additional existing configured NICs for UI access.

See [“Klish menu options for Resiliency Manager”](#) on page 586.

A pre-upgrade checklist is integrated with the `prepare-for-update` command to validate the virtual appliance. This checklist is applicable to Resiliency Manager only. For more details, refer See [“Pre-upgrade checklist”](#) on page 641.



The Resiliency Platform upgrade process is divided into 3 main parts. The upgrade process is applicable for the following platforms:

- Upgrading Resiliency Platform in AWS environment
- Upgrading Resiliency Platform in Azure environment
- Upgrading Resiliency Platform in VMware environment
- Upgrading Resiliency Platform in Hyper-V environment

The detail steps are mentioned according to the respective topics based on the environments. The high level overview of the process for applying updates to the virtual appliances of the Resiliency Platform in version 10.0 is mentioned below:

1. Download and install the required hotfix from Download Center to prepare the virtual appliances for upgrade.
2. Apply the hotfixes to the virtual appliances using KLIST commands to prepare the virtual appliances for upgrade.

See [“Klist menu options for Resiliency Manager”](#) on page 586.

See [“Klish menu options for IMS”](#) on page 598.

3. Detach the data disk of the existing virtual appliance and attach the data disk of the virtual appliance to the new 10.0 appliance.
4. Boot the virtual appliance to start the automatic bootstrap process based on the environment you want to upgrade.
5. In case of Replication Gateway upgrade, you need to replace the existing Replication Gateway with the newer version Replication Gateway.

Before applying updates, ensure that you have performed pre-upgrade tasks for specific scenarios:

See [“Pre-upgrade tasks”](#) on page 636.

The following table consists of the topics you need to refer to the steps mentioned above for applying updates in Veritas Resiliency Platform:

Table 2-177 Applying updates to Resiliency Platform

Step	Task	Steps to perform the task
1	<ol style="list-style-type: none"> a. Download the hotfix from Download Center. b. Apply the hotfix. 	<ol style="list-style-type: none"> a. See “Step 1: Downloading the Resiliency Platform update” on page 636. b.
2	<p>Prepare for update using Klish commands.</p> <p>Prepare the virtual appliances for update using Klish commands.</p>	<p>See “Step 2: Prepare for upgrade” on page 638.</p> <p>See “Klish menu options for Resiliency Manager” on page 586.</p> <p>See “Klish menu options for IMS” on page 598.</p>
3	<p>Detach / attach data disk.</p> <p>You need to perform following steps:</p> <ol style="list-style-type: none"> 1 Detach the data disk from the existing virtual appliance. 2 Create new virtual appliance and attach the existing data disk <p>Choose the environment from the list on which you wish to upgrade the Resiliency Platform.</p>	<p>See “Step 3: Upgrading the Resiliency Platform (Detach / attach the disk)” on page 642.</p>
4	<p>Start Automatic bootstrap process.</p> <p>After the bootstrap process is complete, the upgrade process should automatically start.</p>	<p>See “Step 4: Start the automatic bootstrap process” on page 659.</p>

Table 2-177 Applying updates to Resiliency Platform *(continued)*

Step	Task	Steps to perform the task
5	Replace the Replication Gateway appliance	See “Steps to replace the Replication Gateway appliance” on page 660. For modifying the encryption:
6	Apply update to the single or multiple Resiliency Managers in a domain. In case of multiple Resiliency Managers in the domain, the update needs to be applied on all the Resiliency Managers in synchronization. Note that if you apply update to the Resiliency Manager, then you must apply update to the IMS as well. If you do not update the IMS, the IMS stops data reporting to the Resiliency Manager.	See “Applying update on Resiliency Managers” on page 349.

Apart from upgrading the Resiliency Platform virtual appliances, you need to update following other components too:

Table 2-178 Applying updates to other components of Resiliency Platform

Task	Steps to perform the task
Apply update on the host packages. You must apply update to the Control Host after you apply update to the IMS.	
Apply update on the InfoScale environment In case of InfoScale environment, apply update to the add-on.	
Upgrading the Veritas Replication VIB In case of recovery of VMware virtual machines to on-premises data center using Resiliency Platform Data Mover, If you had configured a resiliency group before applying update to IMS, you need to apply update to the Data Mover bundle.	

Table 2-178 Applying updates to other components of Resiliency Platform
(continued)

Task	Steps to perform the task
Applying updates to the Veritas Data Gateway	

You also have an option of applying a private hotfix, if Veritas support provides you one.

Pre-upgrade tasks

Before upgrading, you may need to perform certain operations in certain scenarios.

Disconnected Resiliency Manager or IMS

If you are using Veritas Resiliency Platform and if any Resiliency Manager or IMS in your resiliency domain is in disconnected state, you must not upgrade. Contact Veritas support to fix the issue before you upgrade the Resiliency Platform components in your resiliency domain.

Rehearsal for recovery of VMware virtual machines to on-premises data center using Resiliency Platform Data Mover

If you have configured resiliency groups for recovery of VMware virtual machines to on-premises data center using Resiliency Platform Data Mover, ensure to perform cleanup rehearsal on the resiliency groups where you have performed rehearsal before applying update on Resiliency Manager.

Before upgrade, if you do not perform cleanup rehearsal on the resiliency groups where you have performed rehearsal, the cleanup rehearsal operation completes successfully after upgrade but other disaster recovery operations such as migrate or recover fails.

See [“About applying updates to Resiliency Platform”](#) on page 631.

Step 1: Downloading the Resiliency Platform update

This topic is a part of the main topic that explains the end-to-end process of applying an update to Veritas Resiliency Platform components. To understand the sequence in which the update needs to be applied to various Veritas Resiliency Platform components, you must see:

Table 2-179 Hotfix details which can applied on

Resiliency Platform version	Hotfix number	Hotfix description	Download location
For version 3.6	3.6.0.20	<p>This is a mandatory hotfix to be applied on Resiliency Platform to upgrade to v10.0. It needs to be installed on all Resiliency Manager, IMS, and Replication Gateways before executing Klish command:.</p> <pre>updates > prepare-for-update</pre> <p>While prepare-for-update is in progress, this hotfix ensures that all the pending processes are complete. It monitors the pending processes for at least 30 minutes. If there are any pending processes in incomplete state prepare-for-update will exit after 30 minutes.</p> <p>For more details refer Pre-upgrade checklist</p>	Hotfix location on Download Center
For version 4.0	4.0.0.6	<p>This is a mandatory hotfix to be applied on Resiliency Platform to upgrade to v10.0. It needs to be installed on all the Resiliency Manager, IMS, and Replication Gateways before executing Klish command <code>updates>prepare-for-update</code>.</p> <p>While prepare-for-update is in progress, this hotfix ensures that all the pending processes are complete. It monitors the pending processes for at least 30 minutes. If there are any pending processes in incomplete state prepare-for-update will exit after 30 minutes.</p> <p>For more details refer Pre-upgrade checklist</p>	Hotfix location on Download Center

To download the update for v3.6, follow below mentioned steps:

- 1 To apply the hotfix, download the .tar.gz files from the above mentioned locations.
- 2 Ensure that you do not unzip the .tar.gz files before uploading it to the virtual appliance. Copy the downloaded files to the required location of the appliance and apply the hotfix.

To download the update for v4.0, follow below mentioned steps:

- 1 To apply the hotfix, download the .hotfix files from the above mentioned locations.
- 2 Ensure that you do not change .hotfix files before uploading it to the virtual appliance. Copy the downloaded files to the required location of the appliance and apply the hotfix.

Next step is to **Prepare for upgrade**, where it is required to prepare the virtual appliance for upgrade.

More Information

Step 2: See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 3: See [“Step 3: Upgrading the Resiliency Platform \(Detach / attach the disk\)”](#) on page 642.

Step 4: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Step 2: Prepare for upgrade

The upgrade process is divided into 3 major parts for all the virtual appliances in which the `Prepare for upgrade` is the first step. You need to perform below steps to prepare the virtual appliances for upgrade using the klish commands:

Prerequisite:

Make sure you have downloaded the appropriate hotfix before upgrading Resiliency Platform to the latest version.

Table 2-180 Hotfix details which can applied on Resiliency Platform

Resiliency Platform version	Hotfix number	Hotfix description	Download location
For version 3.6	3.6.0.20	<p>This is a mandatory hotfix to be applied on Resiliency Platform to upgrade to v10.0. It needs to be installed on all Resiliency Manager, IMS, and Replication Gateways before executing Klish command:</p> <pre>updates> prepare-for-update.</pre> <p>While prepare-for-update is in progress, this hotfix ensures that all the pending processes are complete. It monitors the pending processes for at least 30 minutes. If there are any pending processes in incomplete state prepare-for-update will exit after 30 minutes.</p> <p>For more details refer Pre-upgrade checklist</p>	Hotfix location on Download Center
For version 4.0	4.0.0.6	<p>This is a mandatory hotfix to be applied on Resiliency Platform to upgrade to v10.0. It needs to be installed on all the Resiliency Manager, IMS, and Replication Gateways before executing Klish command</p> <pre>updates> prepare-for-update.</pre> <p>While prepare-for-update is in progress, this hotfix ensures that all the pending processes are complete. It monitors the pending processes for at least 30 minutes. If there are any pending processes in incomplete state prepare-for-update will exit after 30 minutes.</p> <p>For more details refer Pre-upgrade checklist</p>	Hotfix location on Download Center

Steps to perform upgrade

- 1 Login to KLISH menu of the RM appliance and apply the desired hotfix using below command.

```
hotfix > apply-hotfix
```

For example: `hotfix > apply-hotfix <hotfix_number>`

Note: If your data center consists of multiple Resiliency Managers, you have to apply the hotfix on all the RMs.

- 2 On KLISH menu of the RM appliance, execute the below command to prepare the RM for upgrade.

```
updates > prepare-for-update
```

A pre-upgrade checklist is integrated with the `prepare-for-update` command to validate the virtual appliance. This checklist is applicable to Resiliency Manager only. For more details, refer See [“Pre-upgrade checklist”](#) on page 641.

Note: If your data center consists of multiple Resiliency Managers, you have to execute this command on any one RM appliance. This command will prepare all the other RMs for the upgrade. It will also take the backup of all the configurations and power off all the RMs.

- 3 Wait until the RM appliance is powered off automatically. On the RM appliance console, a message appears as “The system is going down for power-off now.

To upgrade the IMS appliance to the latest version, repeat the steps 1-3 on all the available IMS appliances.

The next step is **Upgrading the Resiliency Platform (Attach / detach the disk)** where it is required to detach the data disk of the previous version virtual appliances and attach it to the current version virtual appliance.

See [“Step 3: Upgrading the Resiliency Platform \(Detach / attach the disk\)”](#) on page 642.

Since the upgrade procedure is different for the environments supported by the Resiliency Platform, the steps are different for each environment.

More Information

- Step 1:** See [“Step 1: Downloading the Resiliency Platform update”](#) on page 636.

Step 3: See “[Step 3: Upgrading the Resiliency Platform \(Detach / attach the disk\)](#)” on page 642.

Step 4: See “[Step 4: Start the automatic bootstrap process](#)” on page 659.

Apply updates [About applying updates to Resiliency Platform](#)

Klish menu for Resiliency Manager [Klish menu options for Resiliency Manager](#)

Klish menu for IMS [Klish menu options for IMS](#)

Pre-upgrade checklist

To make upgrade experience smoother a pre-update validation checklist is added as part of prepare-for-update to check and find any discrepancy before starting the upgrade process. This checklist is applicable to Resiliency Manager only.

Following checks are done in pre-upgrade checklist:

1. **Verify pre-upgrade configuration:**

■ **Duplicate objects:**

If the configuration consists of duplicate objects, the post upgrade operations might fail. It checks the database and fetches if any duplicate objects are found. If duplicate objects are found, a warning is displayed.

If the duplicate objects persist, the post-upgrade operations fail. It checks the database and fetches if any duplicate objects are found. If duplicate objects are found, this check will fail and you need to clean the duplicate object with help of Veritas Support and retry.

■ **Stale objects:** These objects might be present in the environment due to any of the following reasons:

■ If any stale Consistency Groups present.

■ If there is any failed rehearsal operation present.

■ If you have performed “force cleanup” which might have left residue. Stale objects block the replace Replication Gateway as part of upgrade activity. To reduce this downtime, these are validated as part of pre-upgrade checks. These stale objects do not block the upgrade process, but it is recommended to clean these duplicate objects with help of Veritas Support. This can be done either before upgrade process or just after upgrading the Resiliency Manager.

2. **Check for on-going activities:** While upgrade is in progress, number of on-going activities is displayed. These activities might get terminated if you confirm to proceed with the upgrade.

3. **Check for on-going reports:** While upgrade is in progress, number of on-going reports is displayed. These activities might get terminated if you confirm to proceed with the upgrade.
4. **Check pending process:** While upgrade is in progress, multiple number of process are pending for execution. You need to make sure these processes are completed or the upgrade operation may exit. While upgrade is in progress, make sure all the processes are complete. If multiple number of process are pending for execution, prepare-for-update will exit after 30 minutes.

See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 3: Upgrading the Resiliency Platform (Detach / attach the disk)

The upgrade process is divided into 3 major parts for all the virtual appliances in which the detach / attach data disk is the second step. In the below mentioned environments, you can detach the data disk of the previous version virtual appliance and attach it to the current / latest version virtual appliance.

- [Upgrading Resiliency Platform in VMware environment](#)
- [Upgrading Resiliency Platform in Hyper-V environment](#)
- [Upgrading Resiliency Platform in AWS environment](#)
- [Upgrading Resiliency Platform in Azure environment](#)

Next step is to **Start the automatic bootstrap process** for the virtual appliances.

More Information:

Step 1: See [“Step 1: Downloading the Resiliency Platform update”](#) on page 636.

Step 2: See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 4: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates: See [“About applying updates to Resiliency Platform”](#) on page 631.

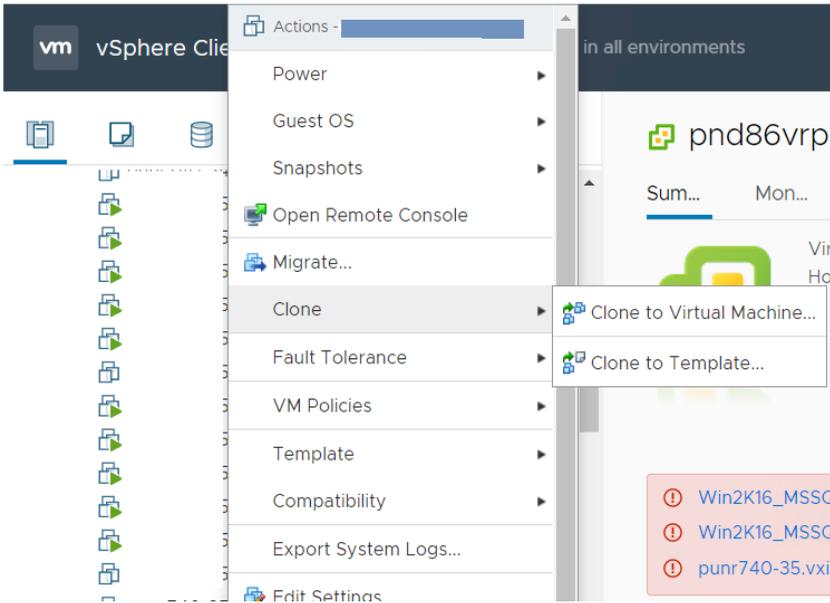
Upgrading Resiliency Platform in VMware environment

The upgrade process is divided into 3 major parts for all the virtual appliances in which the **Attach / Detach data disk** is the second step. You need to perform the below steps to detach the data disk from the previous version virtual appliance and attach it to the current version virtual appliance in the vCenter server:

To attach / detach the data disk

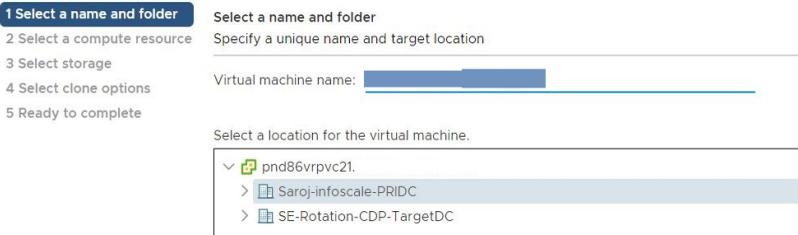
- 1** Ensure that the prepare for upgrade is performed on the virtual appliance. Refer [Step 2: Prepare for upgrade](#).
- 2** Login to vSphere client using admin user credentials.
- 3** Verify that the virtual appliance on which you have applied the updates are shut down.

- 4 To clone the virtual appliance of the previous version perform the below steps:
 - a. Right click on the virtual appliance and select the **Clone > Clone to Virtual Machine** option.



- b. Provide user friendly Virtual machine name to the cloned virtual machine.

- Clone Existing Virtual Machi...



- c. Click **Next**.
 - d. Select the compute resource and click **Next**.
 - e. Select the storage and click **Next**. Complete the clone virtual machine process.

punr [redacted] - Clone Existing Virtual Machi...

1 Select a name and folder
 2 Select a compute resource
 3 Select storage
 4 Select clone options
 5 Ready to complete

Select storage
 Select the storage for the configuration and disk files

Select virtual disk format: Same format as source
 VM Storage Policy: Keep existing VM storage policies

Configure per disk

Name	Capacity	Provisioned	Free	Type
PrimDS1	4.36 TB	5.15 TB	2.14 TB	VM <input type="button" value="v"/>

- 5 Perform the cloning steps on the another virtual appliance.
- 6 Deploy the current version virtual appliance. Do not complete the bootstrap process. To deploy the virtual appliance perform the steps 1-10 from the topic, refer [Deploying the virtual appliance through VMware vSphere Client](#)

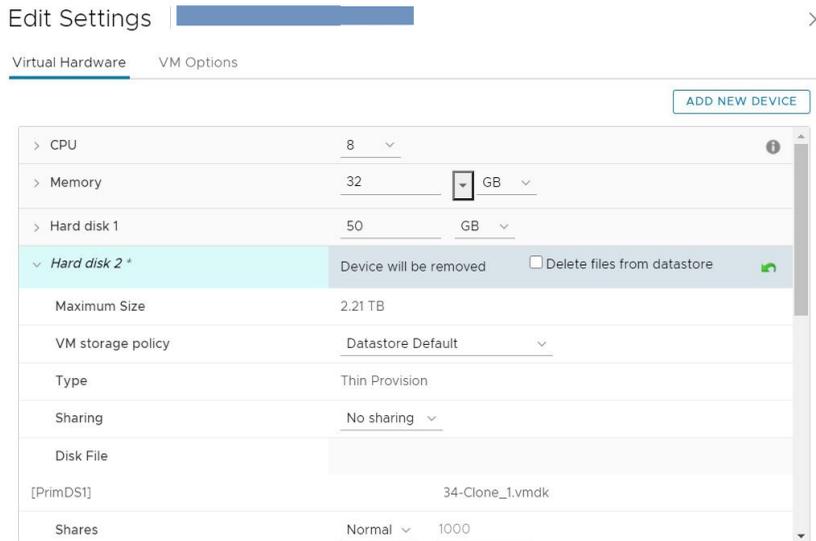
7 The next step is to detach the data disk of the previous version from the appliance. Perform the following steps while you are logged into the vCenter server:

- a. Right click on the appliance and select **Edit settings**.
- b. Select the data disks and choose 'X' to detach the disk.

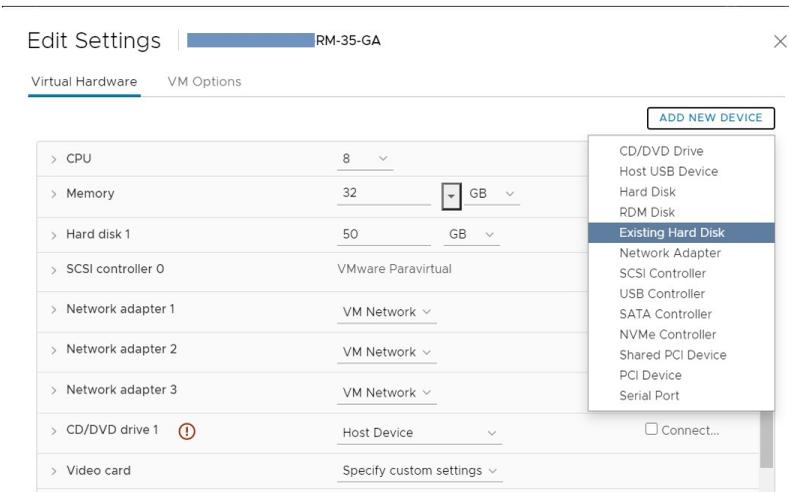
Note: Do not select the 'Delete files from datastore' check box ". This option will delete the selected disk.

c. From the **Disk File** section, note the path of the data disks which is detached in step 7b.

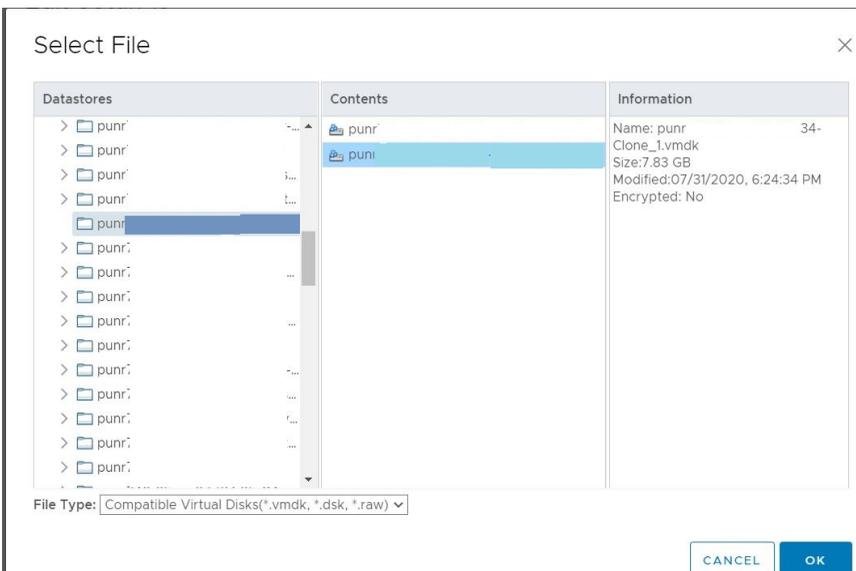
Perform these steps on the other virtual appliances.



- 8 Next step is to attach the data disk to the current version (4.0) virtual appliance which is deployed in step 6 above. Perform the following steps:
 - a. Right click on the virtual appliance and select **Edit settings**.
 - b. Click on the **Add New Device** button and select **Existing Hard Disk** option.



- c. Select the datastore and then select the data disk which is noted in step 7c. Click **OK**.



Perform the steps 8a – 8c on the another virtual appliance to attach the data disk.

- 9 The last step is to power on the virtual appliance. Perform the below steps:
Right click on the appliance and click the **Power On** option. You need to power on the other appliances too.

Note: Attaching the data disk from one virtual appliance to another is supported only during the upgrade process, i.e. after preparing the virtual appliance for upgrade step ([Step 2: Prepare for upgrade](#)). If the data disk of an appliance is changed during normal functioning, it may impact the DR operations.

More Information

Previous Step: See [“Step 2: Prepare for upgrade”](#) on page 638.

Next Step: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

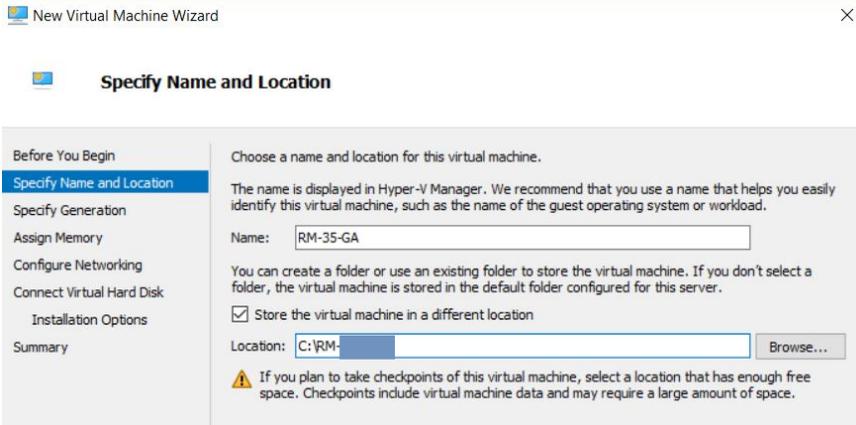
Upgrading Resiliency Platform in Hyper-V environment

The upgrade process is divided into 3 major parts for all the virtual appliances in which the **Create new virtual appliance and attach the data disk** is the second step. You need to perform below steps to create a new appliance and attach the previous version virtual appliance data disk to the current version virtual appliance in the Hyper-V server:

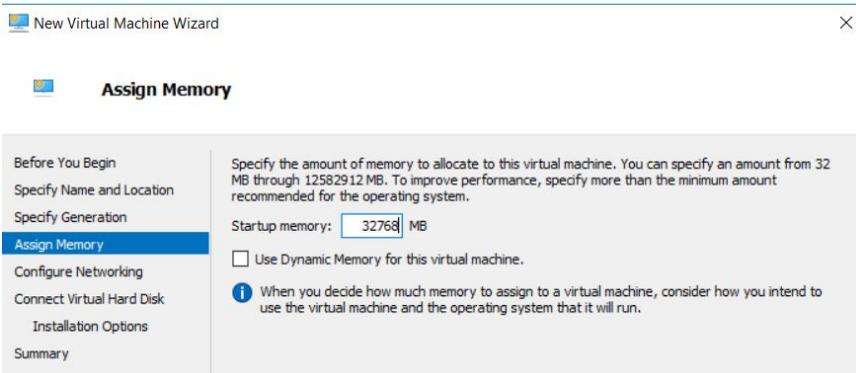
Create new virtual appliance and attach the existing data disk

- 1 Ensure that the prepare for upgrade is performed on the virtual appliance.
Refer [Step 2: Prepare for upgrade](#)
- 2 Login to Hyper-V Manager client using admin user credentials.

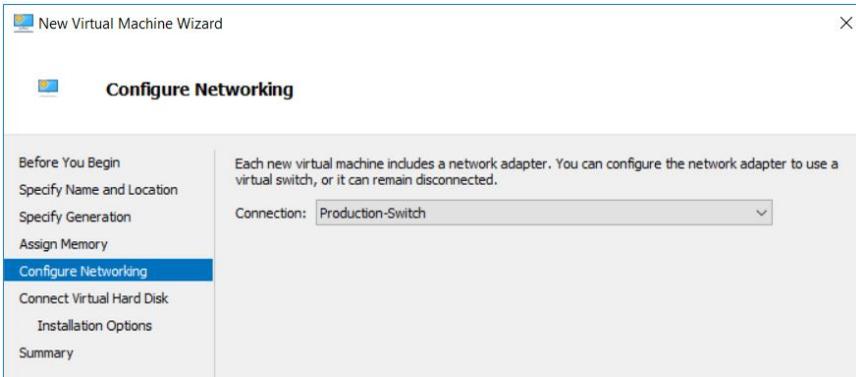
- 3 To deploy the current version virtual appliance, perform the following steps:
 - a. Right-click on the virtual machine and select **New > Virtual Machine**. This opens up the **New Virtual Machine Wizard**.
 - b. Provide the name and location for the virtual machine on the **New Virtual Machine Wizard** and click **Next**.



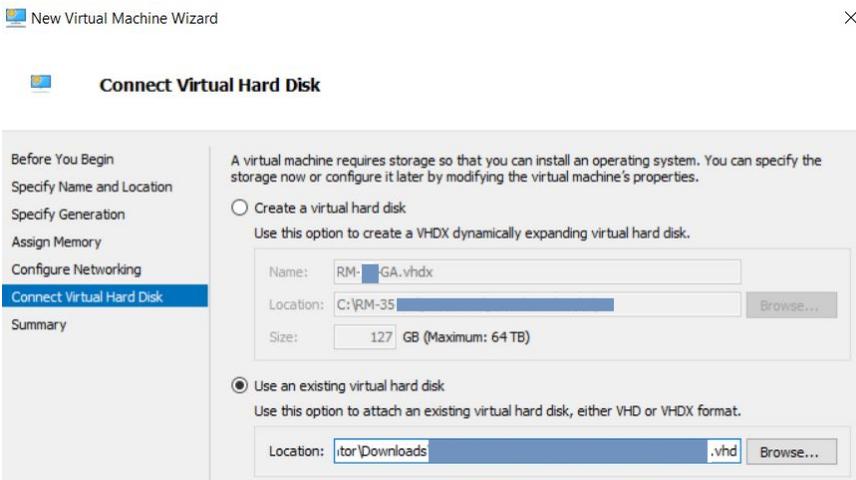
- c. Select the **Generation 1** and click **Next**.
 - d. Provide **Startup memory** in digits based on the system requirement of the Resiliency Platform, refer and click **Next**.



- e. Provide **Configure Networking** as **Production-Switch** and click **Next**.

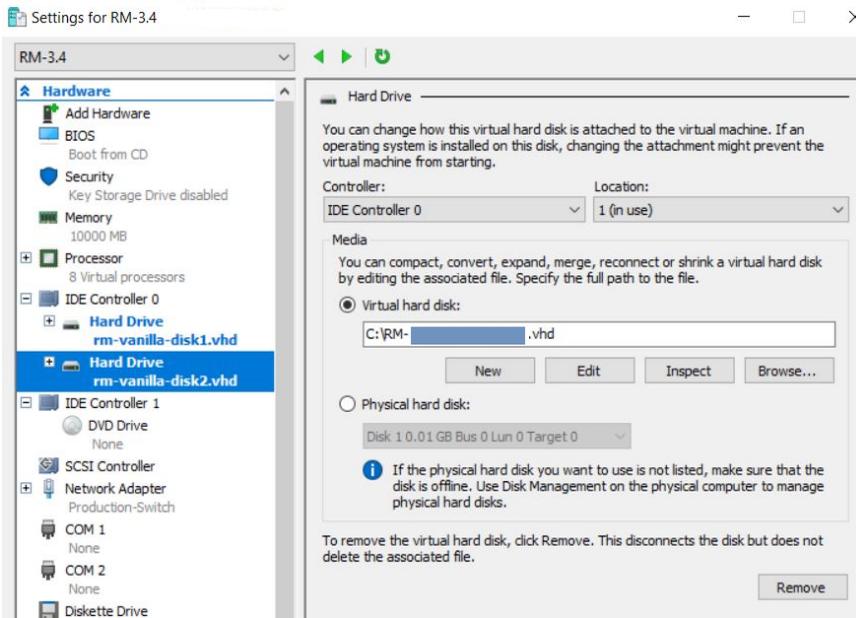


f. In **Connect Virtual Hard Disk** wizard, select **Use an existing virtual hard disk** option and browse to the location where you have saved the version 4.0 hard disk and select Hard Disk 1. Click **Finish**.

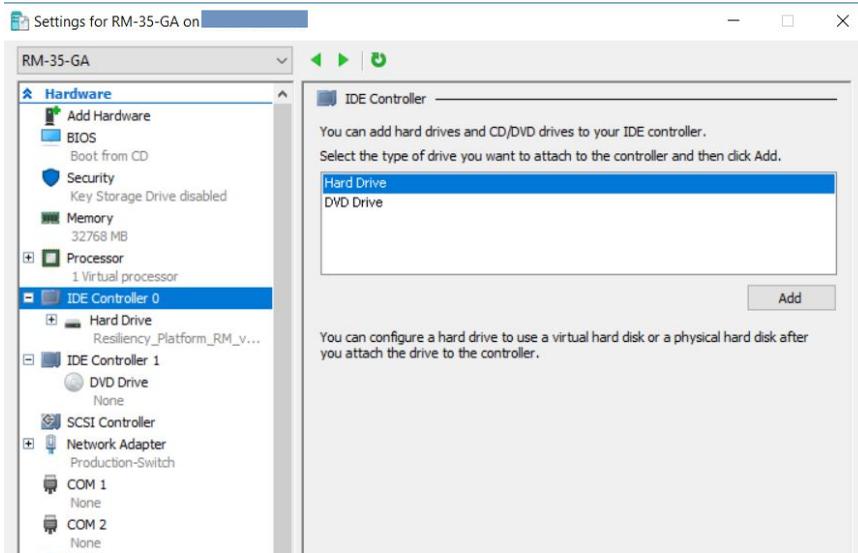


4 Right click on the previous version (3.5 or 3.6) virtual appliance and select **Settings >> IDE Controller >> Hard Disk**. On the **Virtual hard disk** section, select the **Browse** to locate the existing data disk.

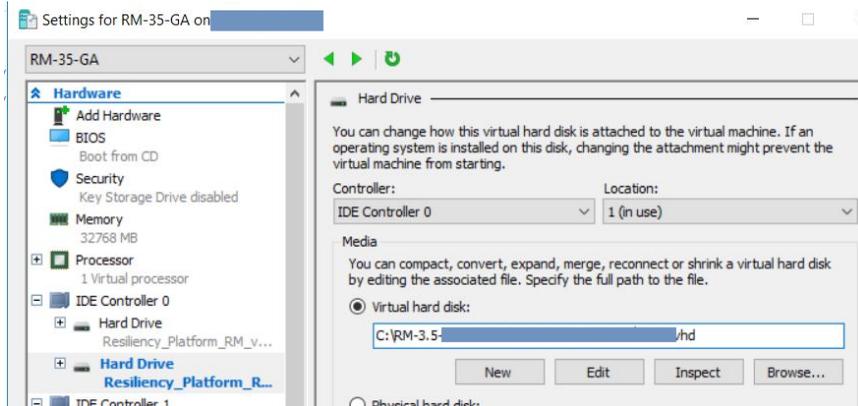
You need to copy the previous (3.5 or 3.6) data disk to the new 4.0 virtual appliance and note the new path of the data disk,



- 6 a. Go to the **Settings** of the 4.0 virtual appliance, select IDE Controller 0, select **Hard Drive** and click **Add**.



- b. Browse the copied data disk location noted in Step 4.



- 7 Click **Apply** after it is enabled and click **OK**.
- 8 Start the virtual appliance by right clicking on the appliance and click **Start** option.

Note: Attaching the data disk from one virtual appliance to another is supported only during the upgrade process, i.e. after preparing the virtual appliance for upgrade step ([Step 2: Prepare for upgrade](#)). If the data disk of an appliance is changed during normal functioning, it may impact the DR operations.

More Information

Previous Step: See [“Step 2: Prepare for upgrade”](#) on page 638.

Next Step: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

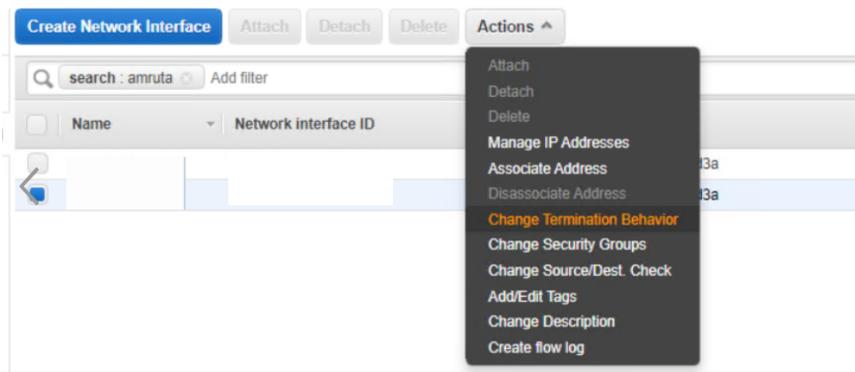
Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Upgrading Resiliency Platform in AWS environment

The upgrade process is divided into 3 major parts for all the virtual appliances in which the **Attach / Detach data disk** is the second step. You need to perform below steps to detach the data disk from previous version virtual appliance and attach it to the version current virtual appliance.

Detach / attach data disk

- 1 Login to AWS portal along with the admin user credentials.
- 2 Navigate to EC2 service and click **Network Interfaces** on the left pane, search for the desired IP address. Provide the name to the NIC and make note of the private IP addresses of the appliance.
- 3 To preserve the NICs of the corresponding IP addresses, navigate to:
 - a. **Actions** drop down.
 - b. Select the **Change Termination Behavior** option.
 - c. Uncheck the **Delete on termination** checkbox and click **Save**.

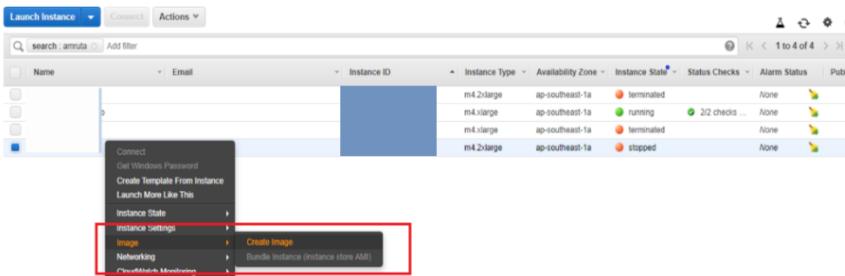


- 4 Make sure you add port 7000 for Resiliency Manager in existing Security Group.
- 5 Navigate to the **Instances** on the left pane. To create an AMI image of the instance i.e backup of the existing version image, navigate to **Actions > Image > Create Image**.

Provide the following details on the **Create Image** wizard and click on the **Create Image**.

Enter Image name:

Enter Image description:



We need to perform this step to roll back the version AMI image if something goes wrong during the upgrade process. For rollback operation, refer the topic See [“Rollback steps in AWS environment”](#) on page 662. .Wait for the Image creation to complete before going to the next page.

- 6 Detach all data disks from the previous version appliance and recommended to tag them with a proper name.
- 7 Delete the older (previous) version virtual appliance to reuse the private IP addresses while configuring the new appliances. Perform the following steps:
 - a. Right-click on older version virtual appliance, select the **Instance State > Terminate**.
 - b. Click on **Yes, Terminate** on the **Terminate Instances** wizard.

- 8 Go to the **AWS Marketplace** and locate **Veritas Resiliency Platform** product. Choose the option of **Upgrade** to upgrade the appropriate virtual appliance.

Parameters
Parameters are defined in your template and allow you to input custom values when you create or update a stack.

Resiliency Manager EC2 Instance Configuration

Instance Name
Enter a name for the new Resiliency Manager EC2 instance

EC2 Instance Type
Select a valid AWS instance type from the list for the new Resiliency Manager instance, or leave as is to automatically select the recommended instance type based on the appliance and region

Key Pair for SSH access
Select an existing EC2 key pair from the list that will be used to enable SSH access to the new Resiliency Manager EC2 instance

Data Volumes from Previous Version
Select ALL data volumes that were attached to the older Resiliency Manager instance. If any volume is missed it would cause issues in the upgrade process.

Network Configuration

VPC ID
Select the VPC in which the new Resiliency Manager EC2 instance will be deployed

Network interface to attach as eth0
Enter the Network interface id of the eth0 network interface, if applicable

Network interface to attach as eth1
Enter the Network interface id of the eth1 network interface, if applicable

Cancel Previous Next

- 9 Provide the required inputs:
 - Instance Name:** Provide appropriate name for the new version instance.
 - EC2 Instance Type:** Select the appropriate instance type to be used.
 - Key Pair for SSH access:** Select the appropriate key pair for the EC2 instance.
 - Data Volumes from Previous Version:** Select all data volumes that were detached from previous version virtual appliance from the drop down list.

Under **Network Configuration** provide the below details:

 - VPC ID:** Provide the VPC id of the previous version virtual appliance.
 - Network interface to attach as eth0:** Enter NIC ids that were noted from previous virtual appliance.
 - Network interface to attach as eth1:** Enter NIC ids that were noted from previous virtual appliance.
- 10 The new version virtual appliance is deployed. Using SSH, continue to complete the upgrade steps.

Note: Attaching the data disk from one virtual appliance to another is supported only during the upgrade process, i.e. after preparing the virtual appliance for upgrade step ([Step 2: Prepare for upgrade](#)). If the data disk of an appliance is changed during normal functioning, it may impact the DR operations.

More Information

Previous Step: See [“Step 2: Prepare for upgrade”](#) on page 638.

Next Step: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Upgrading Resiliency Platform in Azure environment

The upgrade process is divided into 3 major parts for all the virtual appliances in which the **Attach / Detach data disk** is the second step. You need to perform below steps to detach the data disk from the previous version virtual appliance and attach it to the current version virtual appliance.

Detach / attach disk

1 Login to Azure console and navigate to the appliance. Click on the **Capture** option.



- 2** Fill the required details and create an image. This image can be used for the rollback in case the upgrade process fails for any reason.
- 3** Before deleting previous version appliances, note the hostname of the appliance. Use the same hostname while deploying the upgrade offerings.
- 4** Delete the previous version virtual appliance. Ensure that the NICs and volumes of the deleted appliances are not deleted along with the virtual appliance.

- Go to the **Azure Marketplace** and deploy the offering **Veritas Resiliency Platform** by providing following values:

Basic settings for Azure deployment

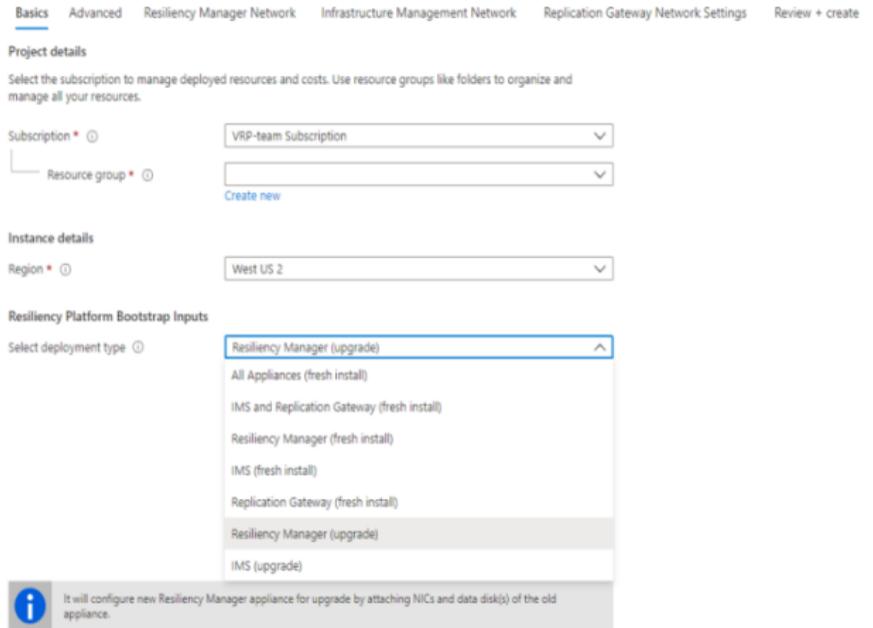


Table: Basic settings for Azure deployment

Input field	Description
Select Deployment Type	Select the appropriate deployment type from the given options. Select 'Resiliency Manager (upgrade)' or 'IMS (upgrade)' for upgrading Resiliency Manager and IMS respectively.
Password for admin user	Set the password for admin user. The admin user and password is later used for configuring the appliances.
Confirm password	Provide same password for confirmation.
Subscription	Select the subscription of Azure account, to be used for deploying the virtual appliances.
Resource group	Specify the name of an existing resource group that contains disks and NICs for the existing appliance, which is being upgraded.

Advance settings for Azure deployment

Basics **Advanced** Resiliency Manager Network Infrastructure Management Network Replication Gateway Network Settings Review + create

Resiliency Manager Settings

Resiliency Manager Instance Name *

Resiliency Manager Hostname *

Resiliency Manager Instance Size * 1x Standard D8s v3
8 vcpus, 32 GB memory
[Change size](#)

Existing Data Disk Name(s) *

Table: Advance settings for Azure deployment

Input field	Description
Instance Name	Name of newer version appliance.
Hostname	Hostname for the newer appliance. This is used if custom DNS is not configured appropriately. Make sure that the hostname is same as the older appliance which is noted before deleting the previous version appliance.
Instance Size	Size of newer version appliance.
Existing Data Disk Name(s)	Provide comma separated list of data disk(s) in same sequence as the LUN number.

Network settings for Azure deployment

Basics Advanced **Resiliency Manager Network** Infrastructure Management Network Replication Gateway Network Settings Review + create

Total NIC(s) Attached *

ETH0 NIC

Type to start filtering...

- AutoGWWU2
Resource group: autoresourcegroupwu2
- AutoMSWU2
Resource group: autoresourcegroupwu2
- bastion-tb-vrp178
Resource group: bastion-tb
- bastion-tb-vrp652
Resource group: bastion-tb

Table: Network settings for Azure deployment

Input field	Description
Total NIC(s) Attached	Select number of NICs attached to older version appliance.
ETHX NIC	Select the appropriate existing NIC attached to the older appliance.

- 6 Click **OK** and review the summary displayed in the **Summary** section.
- 7 Click **OK** to view and accept the terms and conditions.
- 8 Click **Create** to create the instances:
 - This step creates newer version virtual appliance with existing data disk(s) and NICs. To complete the upgrade login to the console with 'admin' username and provided password and continue with further configuration. Make sure you add port 7000 for Resiliency Manager in existing Security Group.

Note: Attaching the data disk from one virtual appliance to another is supported only during the upgrade process, i.e. after preparing the virtual appliance for upgrade step ([Step 2: Prepare for upgrade](#)). If the data disk of an appliance is changed during normal functioning, it may impact the DR operations.

More Information

Previous Step: See [“Step 2: Prepare for upgrade”](#) on page 638.

Next Step: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Step 4: Start the automatic bootstrap process

The upgrade process is divided into 3 major parts for all the virtual appliances in which the `Prepare for upgrade` is the first step. Detaching the data disk and attaching it to the new virtual appliance is the next step. The last step is to start automatic bootstrap process. Below are the steps:

Start the automatic bootstrap process

- 1 While you are in the vSphere client, launch the web console of the RM appliance.
- 2 On the login prompt, provide the default password. Change the password and repeat the same step on the other RM appliances.

- 3 If the appliance detects a data disk from a previous version, you can see a screen showing the network configuration details. Confirm that these details correspond to the correct appliance and proceed to start the bootstrap process.
- 4 Complete the bootstrap process. The upgrade process should start automatically.
- 5 To upgrade the Resiliency Manager refer to [Applying update on Resiliency Managers](#).

In case of multiple Resiliency Managers in the domain, the update needs to be applied on all the Resiliency Managers in synchronization.

Perform steps 1-5 on all the IMSs as well.

Note: If you apply update to the Resiliency Manager, then you must apply update to the IMS as well. If you do not update the IMS, the IMS stops reporting data to the Resiliency Manager.

Once the RM and IMS virtual appliances are upgraded to latest version, the Replication Gateway needs to be replaced with the new Replication Gateway appliance. Refer See [“Steps to replace the Replication Gateway appliance”](#) on page 660.

For more information on upgrading Resiliency Managers, refer the section called [“Applying update on Resiliency Managers”](#)

More Information:

Step 1: See [“Step 1: Downloading the Resiliency Platform update”](#) on page 636.

Step 2: See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 3: See [“Step 3: Upgrading the Resiliency Platform \(Detach / attach the disk\)”](#) on page 642.

Apply updates: See [“About applying updates to Resiliency Platform”](#) on page 631.

Steps to replace the Replication Gateway appliance

Once the RM and IMS virtual appliances are upgraded to the latest version, the Replication Gateway needs to be replaced with the new Replication Gateway appliance. There are two ways to replace the Replication gateway after the RM and IMS are upgraded.

1. Replace the Replication Gateway with same IP address and hostname
2. Replace the Replication Gateway with different IP address and hostname

For more information related to the prerequisites of replace Replication Gateway operation, refer [Replacing a Replication Gateway from a gateway pair](#)

Below are the steps to replace the existing gateway with a new Replication Gateway:

Replace the Replication Gateway with same IP address and hostname

Note: Make sure you do not replace the Replication Gateway and its peer gateway from the Replication Gateway pair simultaneously.

- 1** Login to vSphere client and switch off the existing version of Replication Gateway from source data center by right clicking on the Replication Gateway and select **Power Off** option.
- 2** Next step is to deploy the new Replication Gateway using the same IP address and hostname. You have to select in which environment you are deploying the Replication gateway and accordingly refer to the following topics:
 - [Deploying the virtual appliance through VMware vSphere Client](#)
 - [Deploying the virtual appliance through Hyper-V Manager](#)
 - [Deploying the virtual appliances in AWS through AWS Marketplace](#)
 - [Deploying the virtual appliances in AWS using OVA files](#)
 - [Deploying the virtual appliances in Azure using PowerShell script](#)
 - [Deploying the virtual appliances in Azure through Azure Marketplace](#)
 - [Deploying the virtual appliances in Google Cloud Platform using OVA files](#)
 - [Deploying the virtual appliances in Google Cloud Platform \(GCP\) through GCP Marketplace](#)
- 3** Login into RM and navigate to **Settings > Infrastructure > Data Mover** card mentioned path to check the status of the Replication Gateway is **Faulty**.
- 4** Add the new Replication Gateway which was deployed in step 2. Perform following steps:
 - a. Click on **+ Add Replication Gateway**.
 - b. Provide the IP address and password. Click **Submit**. The workflow is initiated.
- 5** The status of the Replication Gateway is **Healthy** after the workflow is complete.
- 6** Right click on the Replication Gateway and select **Replace Gateway** option.
- 7** Click **OK** on the **Replace Gateway** panel. The Replace Gateway workflow is initiated.

- 8 After the workflow is complete, navigate to the **Settings > Infrastructure > Data Mover** card.
- 9 Verify that the old Replication Gateway is replaced with the new Replication Gateway.
- 10 Navigate to **Assets > Replication Appliance**. The gateway pair status is in **Connected** state.

Perform the same steps to replace the Replication Gateway on the target data center.

For Klish options, See [“Klish menu options for Replication Gateway”](#) on page 612.

Replace the Replication Gateway with different IP address and hostname

- 1 Login to vSphere client and switch off the existing version of Replication Gateway from source data center by right clicking on the Replication Gateway and select **Power Off** option.
- 2 Next step is to deploy the new Replication Gateway using the different IP address and hostname. You have to select in which environment you are deploying the Replication gateway and accordingly refer to the following topics:
[Deploying the virtual appliance through VMware vSphere Client](#)
- 3 To replace the Replication Gateway, repeat the steps 3-10 from the above procedure **Replace the Replication Gateway with same IP address and hostname**.

While upgrading the Replication Gateways from older versions to latest version, it is recommended to replace the Replication Gateway. As part of this process, Replication Gateway with the older version needs to be powered off and should not be powered-on. Hence one of the recommendation is to perform log-gather of the Replication Gateway on older gateway appliance before replacing it.

More Information

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Rollback steps in AWS environment

The rollback operation returns the image or any action return to the previous state. Hence, you may require rollback steps in case the upgrade operation fails.

Steps for rollback image in AWS environment

- 1 Login into the AWS portal.
- 2 Navigate to the virtual appliance of that version on which the error had occurred. Click on the **Delete** option, so that the NICs are free from the new virtual appliance. As performed in the step 6 of topic See [“Upgrading Resiliency Platform in AWS environment”](#) on page 653., preserve the NICs instead of new ones.
- 3 Launch the new instance using the previous version AMI which was noted in step 4 from the topic See [“Upgrading Resiliency Platform in AWS environment”](#) on page 653.. While launching the instance similar to Step 3 in topic See [“Upgrading Resiliency Platform in AWS environment”](#) on page 653., select the preserved NICs instead of creating the new NICs. Also select or create an appropriate IAM role with required permissions as mentioned in See [“Permissions required for IAM roles for Resiliency Manager, IMS, and Replication Gateway”](#) on page 389..
- 4 After the instance is launched, connect it using the admin user and previous version appliance password / SSH key credentials.
- 5 Login into Klish menu of the virtual appliance and execute the below command:

```
updates > rollback-update
```

The previous version virtual appliance is up and online.

More Information

Step 1: See [“Step 2: Prepare for upgrade”](#) on page 638.

Step 2: See [“Upgrading Resiliency Platform in AWS environment”](#) on page 653.

Step 3: See [“Step 4: Start the automatic bootstrap process”](#) on page 659.

Apply updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Rollback steps in Azure environment

The rollback operation returns the image or any action to the previous state in case the upgrade fails.

Steps for rollback image in Azure environment

- 1 Download the `Veritas_Resiliency_Platform_Azure_Upgrade_Scripts_<version>.zip` from the Veritas Download Center.
- 2 Extract the bundle to obtain the `rollbackTemplate.json` and saved it locally.
- 3 Login to Azure portal.

- 4 Browse to the section **Deploy a custom template** option from Azure portal.
- 5 Select the option **Build your own template in the editor**.
- 6 Select the **Load file** option and load the `rollbackTemplate.json` file.
- 7 While clicking the **Save** button, it display the below template:

Home >

Custom deployment

Deploy from a custom template

Select the subscription to manage deployed resources and costs. Use resource groups like folders to organize and manage all your resources.

Subscription * ⓘ

Resource group * ⓘ
[Create new](#)

Instance details

Region * ⓘ

Admin Password *

Appliance Name

Appliance Hostname ⓘ

Appliance Role ⓘ

Old Appliance Version ⓘ

Appliance Deployment Type ⓘ

Appliance Vm Size ⓘ

Image Name ⓘ

Eth0Nic Name ⓘ

Eth1Nic Name ⓘ

Primary Nic

Below are the details of the attributes:

Table 2-181 Custom deployment template attributes

Attributes	Description
Subscription	Select the Subscription Name/ID where the original version Resiliency Platform virtual appliances are deployed.
Resource Group	This should be the same "Resource group" name that was used by the original Resiliency Platform virtual appliances.
Region	This is purely on customer discretion where they want to use.
Admin Password	This password is used by the Resiliency Platform virtual appliance once the rollback deployment is complete. It can be the same password used by the admin on the original virtual appliance (preferred) or provide a new password.
Appliance Name	This would be the display name for the Resiliency Platform virtual appliance
Appliance Hostname	<p>This would be the FQDN or the exact full name of the Resiliency Platform virtual appliance when it was of the original version.</p> <p>Eg: If the original Resiliency Platform Resiliency Manager name was azure-rm-pr.abc.local with display name as Azure-RM, then in that case, the "Appliance name" field will have the value as "Azure-RM" and the "Appliance hostname" will have the value as "azure-rm-pr.abc.local".</p>
Appliance Role	This can either be "Resiliency Manager" or "Infrastructure Management Server" depending on which type of appliance is planned for rollback.
Old Appliance Version	This would be current_version, for example if the older/original version was 4.0.
Appliance Deployment Type	This would be "Marketplace" if the original Resiliency Platform virtual appliance was deployed via Azure Marketplace itself.

Table 2-181 Custom deployment template attributes (*continued*)

Attributes	Description
Appliance virtual machine size	Ensure to select Standard_D8s_v3 for RM and Standard_F8s for IMS.
Image Name	This will be the image name that was created as a part of the “Prepare for upgrade” process. Refer to <code>Apply updates >> Upgrading Resiliency Platform in Azure environment</code> section in Product documentation.
Eth0Nic Name and Eth1Nic Name	Name of the NIC interfaces used by the older/original virtual appliance.
Primary NIC	This would be the “NIC” primarily used for all communications from the virtual appliance. By default, this is eth0 for all Resiliency Platform virtual appliance. (However, ensure to double check and confirm eth0 is the Primary NIC that was used by the older/original virtual appliance.)

8. Once the deployment process is completed, power on the virtual machine and login with the admin user with the password given in the above template.
9. Once successfully logged in, run the command `updates > rollback-update`.

Note: If the rollback is done in a multiple Resiliency Manager setup with above steps, the database rebuild might take some time and until then the “Database service” will be in STOPPED state. It may take up to 30 minutes for database rebuild / repair operation depending upon the size of the database. Once the process is completed, verify the database services using `manage > services status` as ALL.

10. Finally, login to the Resiliency Manager console to confirm all resiliency group details are reflecting as expected.

Step 1: See “[Step 2: Prepare for upgrade](#)” on page 638.

Step 2: See “[Upgrading Resiliency Platform in Azure environment](#)” on page 656.

Step 3: See “[Step 4: Start the automatic bootstrap process](#)” on page 659.

Apply updates See “[About applying updates to Resiliency Platform](#)” on page 631.

Rollback steps in VMware environment

The rollback operation returns the image or any action return to the previous state. Hence, you may require to rollback the steps in case the upgrade operation fails.

Steps to rollback in VMware environment

- 1 Delete the newer version appliance which you have deployed in step 6 in the topic [Upgrading Resiliency Platform in VMware environment](#)
- 2 Power on the older version appliance.
- 3 Run `update > prepare-for-update` command on the older appliance which reverts back the configuration as before

More Information:

Apply Updates: [About applying updates to Resiliency Platform](#)

Rollback steps in Hyper-V environment

The rollback operation returns the image or any action return to the previous state. Hence, you may also require rollback steps in case the upgrade operation fails.

Steps to rollback in Hyper-V environment

- 1 Delete the newer version appliance which you have deployed in step 3 in the topic See [“Upgrading Resiliency Platform in Hyper-V environment”](#) on page 648.
- 2 Power on the older version appliance.
- 3 Run `update > prepare-for-update` command on the older appliance which reverts back the configuration as before.

More Information

Apply Updates See [“About applying updates to Resiliency Platform”](#) on page 631.

Virtual appliance security features

Veritas Resiliency Platform is shipped and deployed in the form of virtual appliances. Following are the security features of Veritas Resiliency Platform virtual appliances:

See [“Operating system security”](#) on page 668.

See [“Management Security”](#) on page 668.

See [“Network security”](#) on page 668.

See [“Access control security”](#) on page 669.

See [“Physical security”](#) on page 668.

Physical security

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

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.

The Control + Alt + Delete key combination has been disabled to avoid any accidental reboot of the virtual appliance. Exec-shield is enabled to protect the virtual appliance from stack, heap, and integer overflows.

Management Security

Only two users are available on the appliance: admin user and support user. These two user accounts are used to access the appliance based on the requirement.

Only admin login is available for the appliance. The password policy of admin login is modified to prompt the user to change the password on the first login.

See [“Password policies for Resiliency Platform virtual appliance”](#) on page 451.

If the admin user password is lost, you need to contact Veritas support for resetting the admin user password.

On successful completion of the product bootstrap, admin user can only access a limited menu of commands through klish. Besides admin user, support user is also supported in the appliance but remote login of support user is disabled. To access the support user, one need to login as an admin and go through **klish**. An option `support > shell` is provided in the **klish** menu to switch the user to support and access the bash shell of support. After selecting this option, the support user is given superuser privileges. Using this option is not recommended and it should be used only with the assistance of technical support.

Timeout of the bash shells of all users is set to 900 seconds.

Network security

The TCP timestamp responses are disabled in Resiliency Platform virtual appliance. Another network security feature of the appliance is that during the product bootstrap process, only those ports that are used by the product for communication and data transfer, are opened through the firewall and all the other communications are blocked.

Uncommon network protocols such as DCCP, SCTP, RDC, TIPC have been disabled so that any process cannot load them dynamically.

Access control security

Resiliency Platform virtual appliance implements certain access control measures. The umask is set to 0700 across the appliance. The access permissions of some of the files such as home folder of root, the log directory etc. is restricted. All the security and the authorization messages are logged into the appliance.

Recovery to cloud data center

This chapter includes the following topics:

- [Recovering VMware virtual machines to AWS](#)
- [Recovering Hyper-V virtual machines to AWS](#)
- [Recovering virtual machines from VMware to AWS using NetBackup Image Sharing](#)
- [Recovering VMware virtual machines to Azure](#)
- [Recovering Hyper-V virtual machines to Azure](#)
- [Recovering virtual machines from Azure / Azure Stack to Azure / Azure Stack](#)
- [Recovering VMware virtual machines to vCloud Director](#)
- [Recovering Hyper-V virtual machines to vCloud Director](#)
- [Recovering VMware virtual machines to vCloud Director without adding vCenter server](#)
- [Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server](#)
- [Recovering virtual machines from vCloud Director to vCloud Director](#)
- [Recovering VMware virtual machines to Orange Recovery Engine](#)
- [Recovering physical machines to AWS using Resiliency Platform Data Mover](#)
- [Recovering physical machines to vCloud Director using Resiliency Platform Data Mover](#)

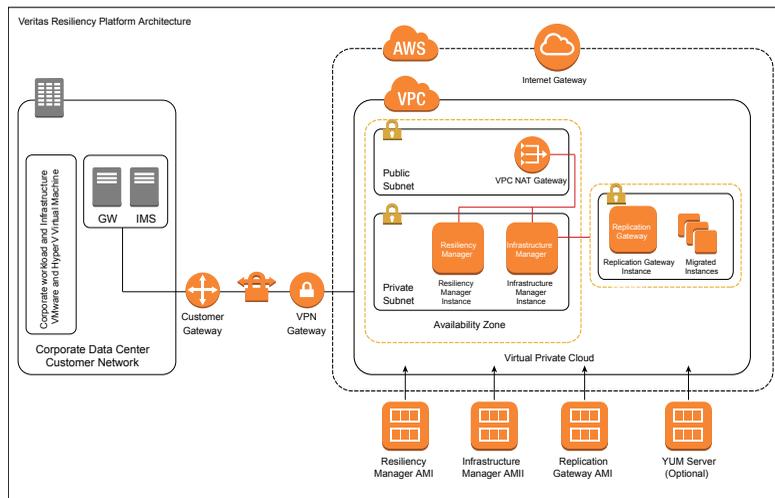
- [Recovering physical machines to Orange Recovery Engine using Resiliency Platform Data Mover](#)
- [Recovering physical machines to Azure using Resiliency Platform Data Mover](#)

Recovering VMware virtual machines to AWS

Using Veritas Resiliency Platform 10.0, you can configure and protect your VMware virtual machines for recovery to AWS using the Resiliency Platform Data Mover.

Instance MetaData Service v2 (IMDSv2) is introduced by AWS as security enhancement over IMDSv1. From version 4.0 of Resiliency Platform, while configuring resiliency group for disaster recovery, the wizard has an option to specify metadata access using **Enforce IMDSv2** option per virtual machine. If this option is true, the virtual machine when migrated to AWS, should use only IMDSv2 mechanism.

Figure 3-1 Overview of deployment Infrastructure for recovery to AWS



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on AWS.

Table 3-1 Recovering VMware virtual machines to AWS

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <p>Overview and Planning Guide</p> <p>Release Notes</p> <p>Checklist for deployment and disaster recovery configuration</p>
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the AWS cloud data center as well as in the premises data center.</p> <p>Refer to the following topics:</p> <p>Download the files required for deployment</p> <p>About deploying the virtual appliances</p> <p>Deploy the Resiliency Platform components in AWS</p> <p>Through AWS marketplace using CloudFormation templates</p> <p>Using OVA files</p> <p>Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center</p> <p>Deploy Data Gateway in AWS environment if you want to use Object Storage for replication</p> <p>Configure the virtual appliances as Veritas Resiliency Platform components</p>

Table 3-1 Recovering VMware virtual machines to AWS (*continued*)

Tasks	More information
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <p>Refer to the following topics:</p> <ul style="list-style-type: none"> Getting started with a new Resiliency Platform configuration <p>Configure the settings for the resiliency domain:</p> <ul style="list-style-type: none"> Adding an IMS Adding a Replication Gateway Adding AWS cloud data center (if not done during getting started wizard) Adding a Data Gateway (only if you want to use Object Storage mode of replication) Managing user authentication and permissions Adding, modifying, or deleting email settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add VMware servers ■ Prepare host for replication
<p>Infrastructure Pairing</p>	<p>For recovering assets to AWS you have to do following infrastructure pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ Create Network group of Cloud Subnets, refer Add network groups (Optional). ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to AWS.

Table 3-1 Recovering VMware virtual machines to AWS (*continued*)

Tasks	More information
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configuring a resiliency group for basic monitoring ■ Prerequisites for configuring resiliency groups for recovery to AWS ■ Configure resiliency groups for recovery to AWS
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Managing virtual business services ■ Managing resiliency plans ■ About evacuation plan
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Performing the rehearsal operation for virtual machines ■ Performing cleanup rehearsal for virtual machines ■ Migrating a resiliency group ■ Recovering resiliency group of virtual machines ■ Performing the resync operation for virtual machines
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ About risks ■ About reports ■ Managing activities

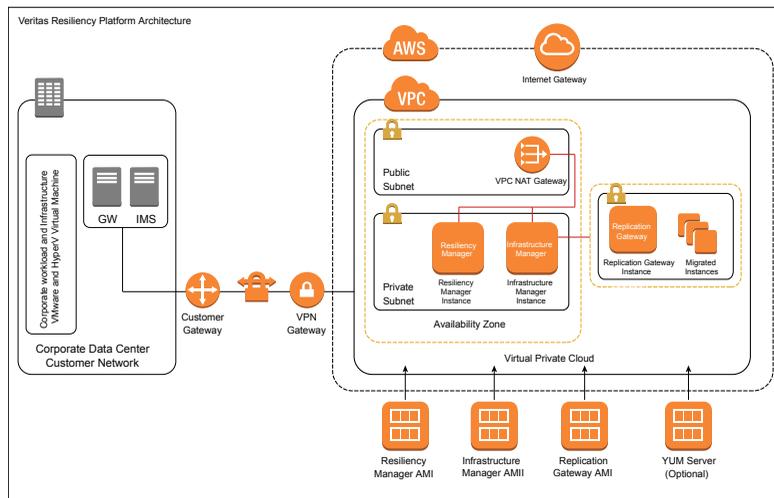
Table 3-1 Recovering VMware virtual machines to AWS (*continued*)

Tasks	More information
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.</p> <ul style="list-style-type: none"> ■ About klish ■ Troubleshoot ■ About applying updates to Resiliency Platform ■ References

Recovering Hyper-V virtual machines to AWS

Using Veritas Resiliency Platform 10.0, you can configure and protect your VMware and Hyper-V virtual machines for recovery to AWS using the Resiliency Platform Data Mover.

Figure 3-2 Overview of deployment Infrastructure for recovery to AWS



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on AWS.

Table 3-2 Recovering Hyper-V virtual machines to AWS

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the AWS cloud data center as well as in the premises data center.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the Resiliency Platform components in AWS by using one of the following methods: <ul style="list-style-type: none"> ■ Through AWS marketplace using CloudFormation templates ■ Using OVA files ■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> ■ Using Hyper-V Manager ■ Deploy Data Gateway in AWS environment if you want to use Object Storage for replication: <ul style="list-style-type: none"> ■ Deploy Data Gateway ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS ■ Configuring Replication Gateways

Table 3-2 Recovering Hyper-V virtual machines to AWS (*continued*)

Tasks	More information
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Add Data Gateway (only if you want to use Object Storage mode of replication) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add Hyper-V servers ■ Prepare host for replication
<p>Infrastructure Pairing</p>	<p>For recovering assets to AWS you have to do following infrastructure pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ Create Network group of Cloud Subnets, refer Add network groups (Optional). ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to AWS.
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Configure resiliency groups for recovery to AWS

Table 3-2 Recovering Hyper-V virtual machines to AWS (*continued*)

Tasks	More information
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ■ Using klish ■ Troubleshooting ■ Updating ■ References

Recovering virtual machines from VMware to AWS using NetBackup Image Sharing

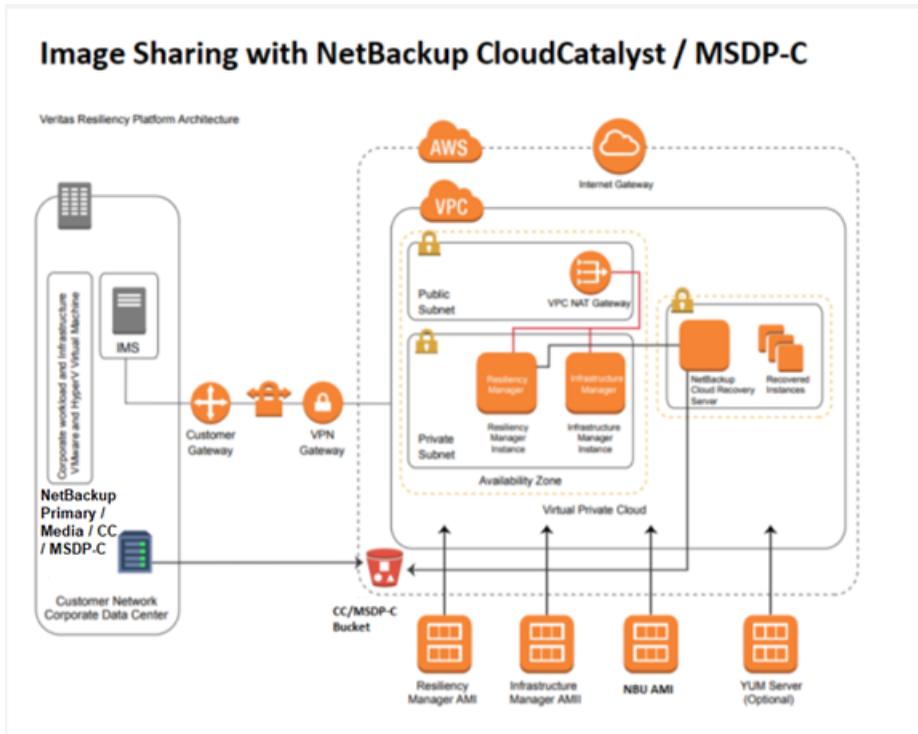
Using the Resiliency Platform, you can recover VMware virtual machine from NetBackup generated backup images that are stored into AWS S3 buckets to the AWS cloud target data center.

In figure the on-premises data center is the source data center and the target data center is an AWS cloud region. The Infrastructure Management Server (IMS) in the on-premises data center discovers the vCenter server and the backup configuration from the NetBackup primary server. NetBackup with the Image Sharing feature, which is available from NetBackup version 8.2 onwards, backs up the virtual machine images along with the image metadata from the on-premise data center into the designated S3 bucket.

For backing up the data in AWS S3 bucket, either NetBackup CloudCatalyst or MSDP-C configuration can be used.

The recovery using these backup images is achieved using a NetBackup Cloud Recovery Server (CRS) virtual appliance that is deployed in the cloud data center. The image metadata stored in the S3 bucket allows the NetBackup CRS to read the image information and create an Amazon Machine Image (AMI). This AMI is then used to provision cloud instances in the cloud data center during the recover operation.

Figure 3-3 Image Sharing with NetBackup CloudCatalyst / MSDP-C



The following table provides the summary of deployment, configuration, and recovery of virtual machines to cloud using NetBackup generated backup images that are stored in S3 bucket using the Image Sharing feature.

Table 3-3 Recovering virtual machines using NetBackup images

Tasks	More information
<p data-bbox="124 1260 360 1286">Plan your environment</p> 	<p data-bbox="413 1260 1216 1373">Refer to the <i>Veritas Resiliency Platform Overview and Planning Guide</i> to know about the product, its components, features, and capabilities. Refer to the <i>Veritas Resiliency Platform Release Notes</i> for release information such as main features, known issues, and limitations.</p> <p data-bbox="413 1390 700 1416">Overview and Planning Guide</p> <p data-bbox="413 1433 556 1459">Release Notes</p> <p data-bbox="413 1477 1180 1529">Ensure that the configuration details in your environment match the requirements mentioned in the checklist.</p> <p data-bbox="413 1546 1190 1572">Checklist for recovery of VMware virtual machines to AWS cloud using NetBackup</p>

Table 3-3 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances for Resiliency Manager and IMS in both the data centers.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment Downloading the Veritas Resiliency Platform virtual appliances ■ About deploying the virtual appliances About deploying the Resiliency Platform virtual appliances ■ Deploy the Resiliency Platform components in AWS by using one of the following methods: <ul style="list-style-type: none"> ■ Deploying the virtual appliances in AWS through AWS Marketplace ■ Deploying the virtual appliances in AWS using OVA files ■ Deploy the virtual appliances for one or more IMS in the premises data center: <ul style="list-style-type: none"> ■ Deploying the virtual appliance through VMware vSphere Client ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ See “About configuring the Resiliency Platform components ” on page 437. About configuring the Resiliency Platform components ■ See “Prerequisites for configuring Resiliency Platform components” on page 438. Prerequisites for configuring Resiliency Platform components ■ Deploy the Cloud Recovery Server (CRS) Deploy CRS <p>For more information on CRS, refer to <i>About Image Sharing using MSDP cloud</i> topic in <i>NetBackup Deduplication Guide</i>.</p>
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Getting started with a new Resiliency Platform configuration ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Adding an IMS ■ Adding AWS cloud data center ■ Managing user authentication and permissions ■ For secure communication, refer Managing security
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Adding VMware virtualization servers ■ Adding NetBackup primary server ■ Adding NetBackup Cloud Recovery Server (CRS)

Table 3-3 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Infrastructure Pairing</p>	<p>For recovering assets to VMware you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ Create Network group of cloud subnets. Adding a network group ■ Customize DNS Configuring DNS server settings for a data center ■ Create network mappings. Network pairs for recovering virtual machines to AWS
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for recovery on local or remote data center.</p> <ul style="list-style-type: none"> ■ Configuring a resiliency group for basic monitoring ■ Managing VMware virtual machines for remote recovery to AWS cloud using NetBackup images
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Managing virtual business services ■ Managing resiliency plans ■ About evacuation plan
<p>Perform recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform recover (local or remote) operations on the resiliency groups.</p> <p>Note: The rehearsal operation is not supported from AWS to VMware.</p> <ul style="list-style-type: none"> ■ Performing the rehearsal operation on virtual machines from VMware to AWS using NetBackup Image Sharing ■ Performing cleanup rehearsal for virtual machines ■ Recovering virtual machines to cloud (AWS) using NetBackup Image Sharing

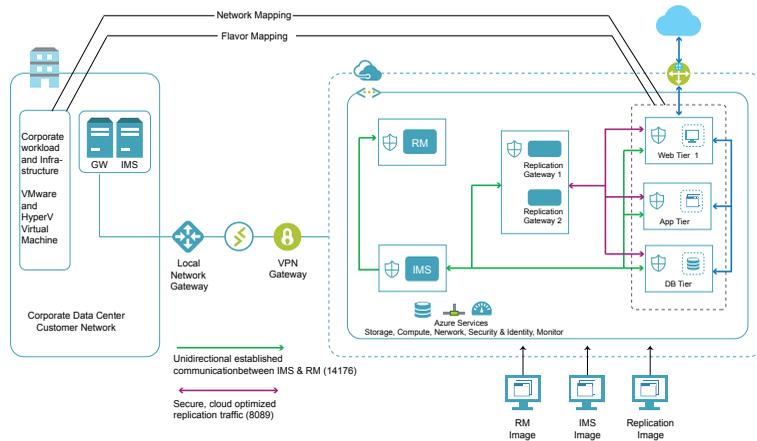
Table 3-3 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ About risks ■ About reports ■ Managing activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ■ About klish ■ Troubleshoot ■ About applying updates to Resiliency Platform ■ References

Recovering VMware virtual machines to Azure

Using Veritas Resiliency Platform 10.0, you can configure and protect your VMware virtual machines for recovery to Azure using the Resiliency Platform Data Mover.

Figure 3-4 Overview of deployment Infrastructure for recovery to Azure



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on Azure.

Table 3-4 Recovering VMware virtual machines to Azure

Tasks	More information
Plan your environment 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 3-4 Recovering VMware virtual machines to Azure (*continued*)

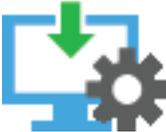
Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the Azure cloud data center as well as in the premises data center.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway in the Azure cloud data center, using any of the following options: <ul style="list-style-type: none"> See “Deploying the virtual appliances in Azure using PowerShell script” on page 396. See “Deploying the virtual appliances in Azure through Azure Marketplace” on page 405. ■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> ■ Using VMware vSphere client ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS ■ Configuring Replication Gateways
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add VMware servers ■ Prepare host for replication

Table 3-4 Recovering VMware virtual machines to Azure (*continued*)

Tasks	More information
Infrastructure Pairing	<p>For recovering assets to Azure you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to Azure.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Configure resiliency groups for recovery to Azure
Advance features 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
Perform remote recovery operations 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync

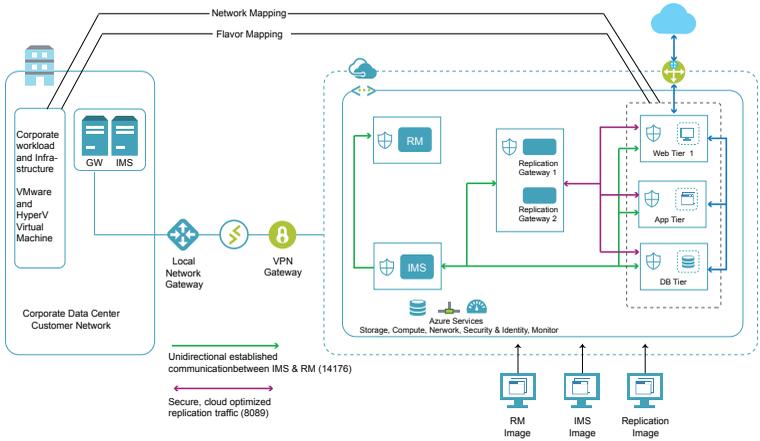
Table 3-4 Recovering VMware virtual machines to Azure (*continued*)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

Recovering Hyper-V virtual machines to Azure

Using Veritas Resiliency Platform 10.0, you can configure and protect your Hyper-V virtual machines for recovery to Azure using the Resiliency Platform Data Mover.

Figure 3-5 Overview of deployment Infrastructure for recovery to Azure



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on Azure.

Table 3-5 Recovering Hyper-V virtual machines to Azure

Tasks	More information
<p data-bbox="126 1017 358 1038">Plan your environment</p> 	<p data-bbox="415 1017 1214 1130">Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul data-bbox="415 1147 1018 1237" style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 3-5 Recovering Hyper-V virtual machines to Azure (*continued*)

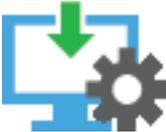
Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the Azure cloud data center as well as in the premises data center.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway in the Azure cloud data center using any of the following options: <ul style="list-style-type: none"> See “Deploying the virtual appliances in Azure using PowerShell script” on page 396. See “Deploying the virtual appliances in Azure through Azure Marketplace” on page 405. ■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> ■ Using Hyper-V Manager ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS ■ Configuring Replication Gateways
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add Hyper-V servers ■ Prepare host for replication

Table 3-5 Recovering Hyper-V virtual machines to Azure (*continued*)

Tasks	More information
Infrastructure Pairing	<p>For recovering assets to Azure you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to Azure.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Configure resiliency groups for recovery to Azure
Advance features 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
Perform remote recovery operations 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync

Table 3-5 Recovering Hyper-V virtual machines to Azure (*continued*)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ■ Using klish ■ Troubleshooting ■ Updating ■ References

Recovering virtual machines from Azure / Azure Stack to Azure / Azure Stack

Using Veritas Resiliency Platform 10.0, you can configure and protect your virtual machines for recovery from Azure to Azure using the Resiliency Platform Data Mover which includes combination of :

- Azure Stack to Azure Stack
- Azure Stack to Azure region
- Azure region to Azure region
- Azure region to Azure Stack

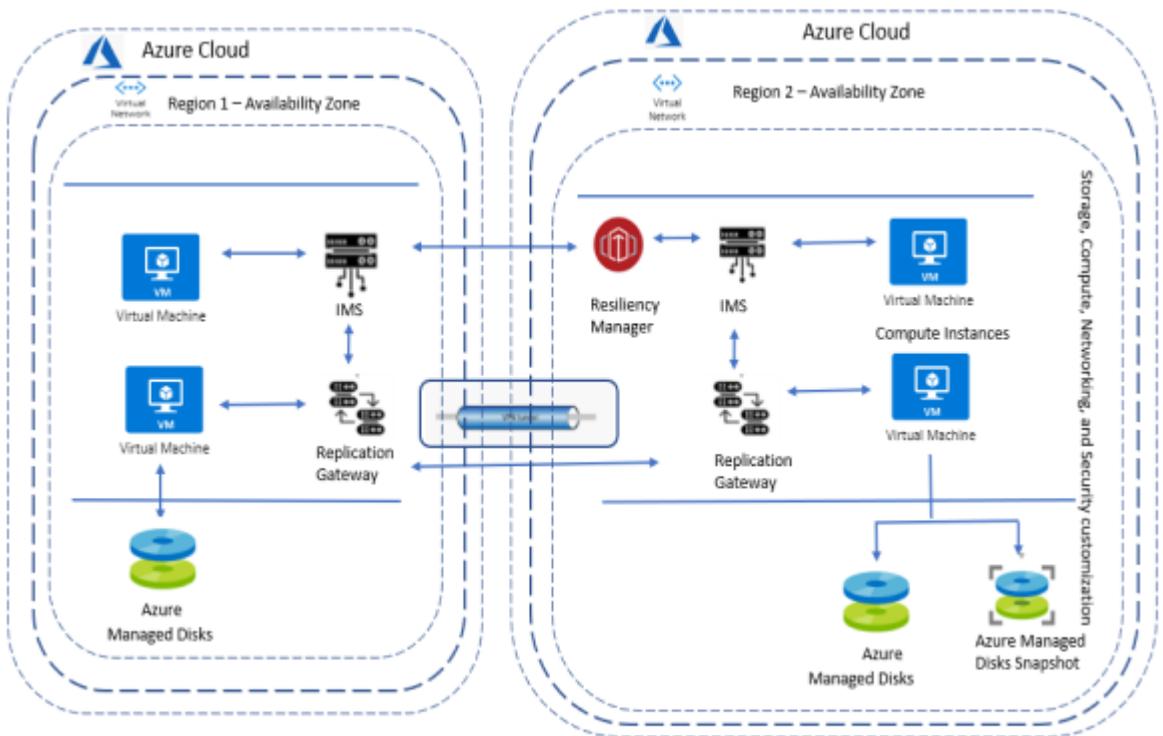
You can use the same or different Azure cloud subscriptions for using the Azure resources. Following are the features supported for this use case:

1. NRT discovery support is enabled for the virtual machines which are present at the target data center. Using NRT, the real-time updates related to network and virtual machines are discovered.

Recovering virtual machines from Azure / Azure Stack to Azure / Azure Stack

2. You can now edit the migrated virtual machine, add/ remove the virtual machine from a resiliency group on the target data center.
3. Rehearsal operation is enabled when target data center is Azure and vice versa.
4. Network customization feature is added where you can enable or disable IP customization on both source and target data center.

Figure 3-6 Overview of deployment Infrastructure for recovery from Azure to Azure



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on Azure to Azure which includes (Azure Stack to Azure Stack, Azure Stack to Azure region, Azure region to Azure region and Azure region to Azure Stack).

Table 3-6 Recovering virtual machines from Azure to Azure

Tasks	More information
<p data-bbox="126 326 360 352">Plan your environment</p> 	<p data-bbox="412 326 1216 439">Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> <li data-bbox="412 458 729 484">■ Overview and Planning Guide <li data-bbox="412 491 588 517">■ Release Notes <li data-bbox="412 524 1018 550">■ Checklist for deployment and disaster recovery configuration
<p data-bbox="126 604 384 661">Deploy and configure the virtual appliances</p> 	<p data-bbox="412 604 1216 661">Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the Azure cloud data center on source and target data centers.</p> <ul style="list-style-type: none"> <li data-bbox="412 680 854 706">■ Download the files required for deployment <li data-bbox="412 713 1216 826">■ Deploy the virtual appliances Infrastructure Management Server (IMS) and Replication Gateway on Azure cloud at both the data centers using any one of the following options. Resiliency Manager should be deployed either on source or on target data center. <ul style="list-style-type: none"> <li data-bbox="444 831 1216 857">See “Deploying the virtual appliances in Azure using PowerShell script” on page 396. <li data-bbox="444 864 1153 921">See “Deploying the virtual appliances in Azure through Azure Marketplace” on page 405. <li data-bbox="444 928 1177 986">See “Deploying the virtual appliances in Azure Stack using PowerShell script” on page 401. <li data-bbox="444 992 1184 1050">See “Deploy virtual appliances in Azure Stack using Azure Stack Marketplace” on page 412. <li data-bbox="412 1057 1161 1083">■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> <li data-bbox="444 1090 852 1116">■ About configuring the virtual appliances <li data-bbox="444 1123 951 1149">■ Prerequisites for configuring the virtual appliances <li data-bbox="444 1156 852 1182">■ Configuring Resiliency Manager or IMS <li data-bbox="444 1189 801 1215">■ Configuring Replication Gateways

Table 3-6 Recovering virtual machines from Azure to Azure (*continued*)

Tasks	More information
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add another cloud data center ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings <p>Using the Resiliency Platform console, you can add one or more Azure Stack private cloud instances to the non-cloud datacenter (premise) at source, or at target or both data centers. Azure Stack private cloud can be added to non-cloud datacenter only.</p> <ul style="list-style-type: none"> ■ Adding Azure Stack private cloud instance ■ For secure communication, refer Managing security
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Prepare host for replication
<p>Infrastructure Pairing</p>	<p>For recovering assets to Azure you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to Azure.

Table 3-6 Recovering virtual machines from Azure to Azure (*continued*)

Tasks	More information
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Configure resiliency groups for recovery from Azure cloud to Azure cloud
<p>Advance features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Evacuation plans
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <p>From version 10.0, Resiliency Platform supports rehearsal operation from Azure region to Azure region along with (Azure Stack to Azure Stack, Azure Stack to Azure region, and Azure region to Azure Stack) .</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities

Table 3-6 Recovering virtual machines from Azure to Azure (*continued*)

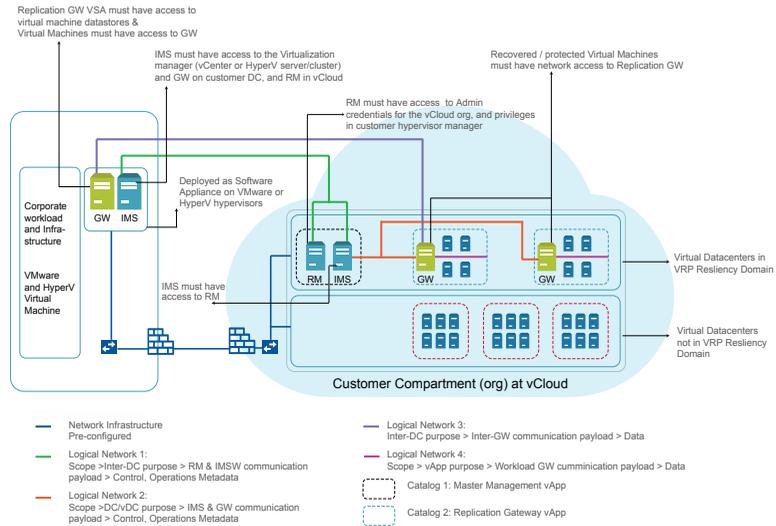
Tasks	More information
Miscellaneous references 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none">■ Using klish■ Troubleshooting■ Updating■ References

Recovering VMware virtual machines to vCloud Director

Using Veritas Resiliency Platform 10.0, you can configure and protect your VMware virtual machines for recovery to vCloud Director using the Resiliency Platform Data Mover.

Before starting the product deployment in your data center, ensure that the cloud tenant is created for you and you have the cloud credentials to access it.

Figure 3-7 Overview of deployment infrastructure for recovery to vCloud Director



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on vCloud Director. These operations can be performed by the end user or the service subscriber.

Table 3-7 Recovering VMware virtual machines to vCloud Director

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment matches the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 3-7 Recovering VMware virtual machines to vCloud Director
(continued)

Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the premises as well as cloud data center.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances in vCloud Director for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway. Each virtual data center in vCloud is represented as an individual data center in Resiliency Platform. If you have multiple virtual data centers, you need to create multiple data centers in Resiliency Platform and then deploy Resiliency Manager and IMS in one virtual data center and only IMS in rest of the virtual data centers: <ul style="list-style-type: none"> ■ Using vCloud Director ■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> ■ Using VMware vSphere client ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS ■ Configuring Replication Gateways
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add VMware servers ■ Prepare host for replication

Table 3-7 Recovering VMware virtual machines to vCloud Director
(continued)

Tasks	More information
<p>Infrastructure Pairing</p>	<p>For recovering assets to vCloud Director you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ Create Network group of vLAN/Port Group, refer Add network groups (Optional). ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to vCloud Director.
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Manage resiliency groups for remote recovery
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
<p>Perform remote recovery operations</p> 	<p>Once you have organized your assets into resiliency groups, you can perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync <p>Note that, Rehearsal and Cleanup Rehearsal operations are not supported for recovery to vCloud Director.</p>

Table 3-7 Recovering VMware virtual machines to vCloud Director
(continued)

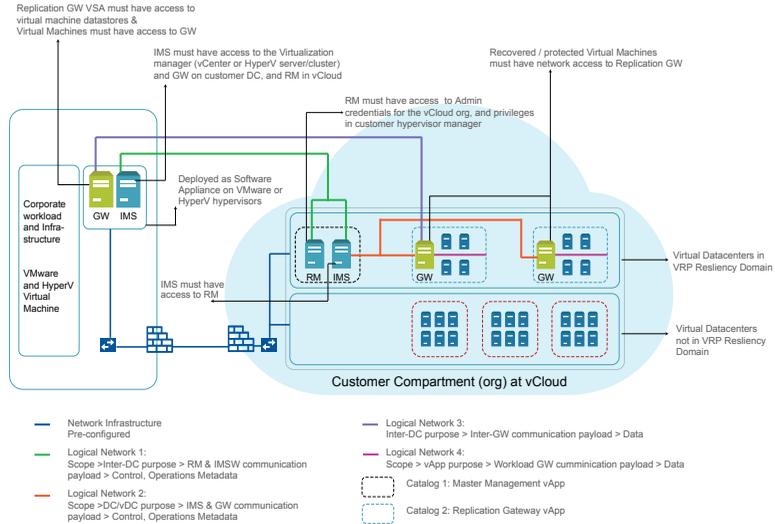
Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

Recovering Hyper-V virtual machines to vCloud Director

Using Veritas Resiliency Platform 10.0, you can configure and protect your Hyper-V virtual machines for recovery to vCloud Director using the Resiliency Platform Data Mover.

Before starting the product deployment in your data center, ensure that the cloud tenant is created for you and you have the cloud credentials to access it.

Figure 3-8 Overview of deployment infrastructure for recovery to vCloud Director



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on vCloud Director. These operations can be performed by the end user or the service subscriber.

Table 3-8 Recovering Hyper-V virtual machines to vCloud Director

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 3-8 Recovering Hyper-V virtual machines to vCloud Director
(continued)

Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the premises as well as cloud data center.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances in vCloud Director for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway. If you have multiple virtual data centers, deploy Resiliency Manager and IMS in one virtual data center and only IMS in rest of the virtual data centers: <ul style="list-style-type: none"> ■ Using vCloud Director ■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> ■ Using Hyper-V Manager ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS ■ Configuring Replication Gateways
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add Hyper-V servers ■ Prepare host for replication

Table 3-8 Recovering Hyper-V virtual machines to vCloud Director
(continued)

Tasks	More information
Infrastructure Pairing	<p>For recovering assets to vCloud Director you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ Create Network group of vLAN/Port Group, refer Add network groups (Optional). ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to vCloud Director.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Manage resiliency groups for remote recovery
Advanced features 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
Perform remote recovery operations 	<p>Once you have organized your assets into resiliency groups, you can perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync <p>Note that, Rehearsal and Cleanup Rehearsal operations are not supported for recovery to vCloud Director.</p>

Table 3-8 Recovering Hyper-V virtual machines to vCloud Director
(continued)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

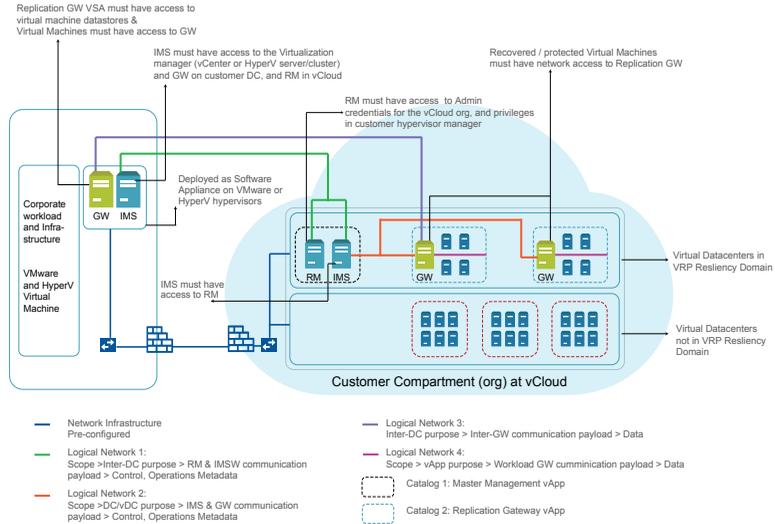
Recovering VMware virtual machines to vCloud Director without adding vCenter server

Using Veritas Resiliency Platform 10.0, you can configure and protect your VMware virtual machines for recovery to vCloud Director using the Resiliency Platform Data Mover without adding the vCenter server.

Before starting the product deployment in your data center, ensure that the cloud tenant is created for you and you have the cloud credentials to access it.

Recovering VMware virtual machines to vCloud Director without adding vCenter server

Figure 3-9 Overview of deployment infrastructure for recovery to vCloud Director



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on vCloud Director. These operations can be performed by the end user or the service subscriber.

Table 3-9 Recovering VMware virtual machines to vCloud Director without adding vCenter server

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 3-9 Recovering VMware virtual machines to vCloud Director without adding vCenter server (*continued*)

Tasks	More information
<p data-bbox="126 354 385 409">Deploy and configure the virtual appliances</p> 	<p data-bbox="413 354 1201 409">Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the premises as well as cloud data center.</p> <ul style="list-style-type: none"> <li data-bbox="413 427 854 453">■ Download the files required for deployment <li data-bbox="413 461 807 487">■ About deploying the virtual appliances <li data-bbox="413 496 1217 609">■ Deploy the virtual appliances in vCloud Director for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway. If you have multiple virtual data centers, deploy Resiliency Manager and IMS in one virtual data center and only IMS in rest of the virtual data centers: <ul style="list-style-type: none"> <li data-bbox="444 618 686 644">■ Using vCloud Director <li data-bbox="413 652 1217 704">■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> <li data-bbox="444 713 760 739">■ Using VMware vSphere client <li data-bbox="413 748 1163 774">■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> <li data-bbox="444 782 852 808">■ About configuring the virtual appliances <li data-bbox="444 817 852 843">■ Configuring Resiliency Manager or IMS <li data-bbox="444 852 801 878">■ Configuring Replication Gateways
<p data-bbox="126 892 337 947">Set up the resiliency domain</p> 	<p data-bbox="413 892 1217 982">Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> <li data-bbox="413 999 982 1025">■ Create the resiliency domain using getting started wizard <li data-bbox="413 1034 895 1060">■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> <li data-bbox="444 1069 559 1095">■ Add IMS <li data-bbox="444 1104 727 1130">■ Add Replication Gateways <li data-bbox="444 1138 1083 1164">■ Add cloud data center (if not done during getting started wizard) <li data-bbox="444 1173 895 1199">■ Manage user authentication and permission <li data-bbox="444 1208 995 1234">■ Manage alerts, notifications, and other product settings
<p data-bbox="126 1248 374 1274">Add asset infrastructure</p> 	<p data-bbox="413 1248 1217 1338">Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> <li data-bbox="413 1355 704 1381">■ Prepare host for replication

Recovering VMware virtual machines to vCloud Director without adding vCenter server

Table 3-9 Recovering VMware virtual machines to vCloud Director without adding vCenter server (*continued*)

Tasks	More information
Infrastructure Pairing	<p>For recovering assets to vCloud Director you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to vCloud Director.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity.</p> <ul style="list-style-type: none"> ■ Manage resiliency groups for remote recovery
Advanced features 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
Perform remote recovery operations 	<p>Once you have organized your assets into resiliency groups, you can perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync <p>Note that, Rehearsal and Cleanup Rehearsal operations are not supported for recovery to vCloud Director.</p>

Table 3-9 Recovering VMware virtual machines to vCloud Director without adding vCenter server (*continued*)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

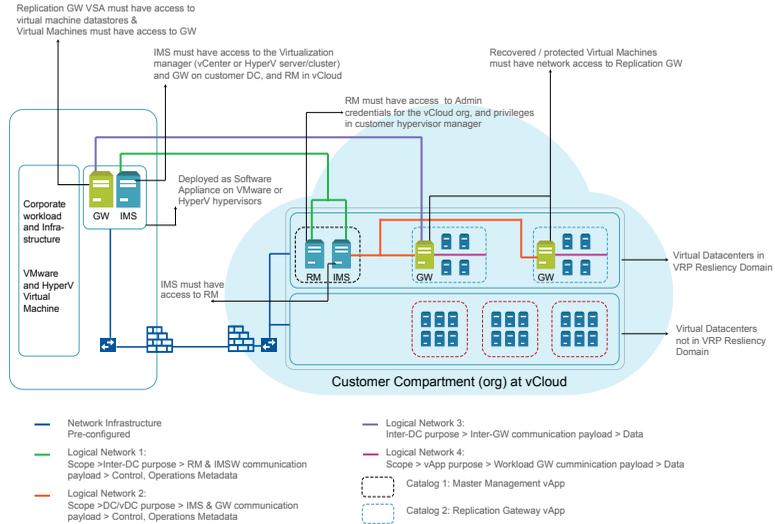
Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server

Using Veritas Resiliency Platform 10.0, you can configure and protect your Hyper-V virtual machines for recovery to vCloud Director using the Resiliency Platform Data Mover without adding Hyper-V server.

Before starting the product deployment in your data center, ensure that the cloud tenant is created for you and you have the cloud credentials to access it.

Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server

Figure 3-10 Overview of deployment infrastructure for recovery to vCloud Director



The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on vCloud Director. These operations can be performed by the end user or the service subscriber.

Table 3-10 Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 3-10 Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server (*continued*)

Tasks	More information
<p data-bbox="126 354 385 409">Deploy and configure the virtual appliances</p> 	<p data-bbox="412 354 1201 409">Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the premises as well as cloud data center.</p> <ul style="list-style-type: none"> <li data-bbox="412 427 852 453">■ Download the files required for deployment <li data-bbox="412 461 805 487">■ About deploying the virtual appliances <li data-bbox="412 496 1217 609">■ Deploy the virtual appliances in vCloud Director for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway. If you have multiple virtual data centers, deploy Resiliency Manager and IMS in one virtual data center and only IMS in rest of the virtual data centers: <ul style="list-style-type: none"> <li data-bbox="444 618 686 644">■ Using vCloud Director <li data-bbox="412 652 1217 704">■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> <li data-bbox="444 713 706 739">■ Using Hyper-V Manager <li data-bbox="412 748 1161 774">■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> <li data-bbox="444 782 852 808">■ About configuring the virtual appliances <li data-bbox="444 817 852 843">■ Configuring Resiliency Manager or IMS <li data-bbox="444 852 801 878">■ Configuring Replication Gateways
<p data-bbox="126 892 337 947">Set up the resiliency domain</p> 	<p data-bbox="412 892 1217 982">Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> <li data-bbox="412 999 982 1025">■ Create the resiliency domain using getting started wizard <li data-bbox="412 1034 899 1060">■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> <li data-bbox="444 1069 563 1095">■ Add IMS <li data-bbox="444 1104 731 1130">■ Add Replication Gateways <li data-bbox="444 1138 1080 1164">■ Add cloud data center (if not done during getting started wizard) <li data-bbox="444 1173 892 1199">■ Manage user authentication and permission <li data-bbox="444 1208 995 1234">■ Manage alerts, notifications, and other product settings
<p data-bbox="126 1248 374 1274">Add asset infrastructure</p> 	<p data-bbox="412 1248 1217 1338">Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> <li data-bbox="412 1355 704 1381">■ Prepare host for replication

Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server

Table 3-10 Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server (*continued*)

Tasks	More information
Infrastructure Pairing	<p>For recovering assets to vCloud Director you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering virtual machines to vCloud Director.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity.</p> <ul style="list-style-type: none"> ■ Manage resiliency groups for remote recovery
Advanced features 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
Perform remote recovery operations 	<p>Once you have organized your assets into resiliency groups, you can perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync <p>Note that, Rehearsal and Cleanup Rehearsal operations are not supported for recovery to vCloud Director.</p>

Table 3-10 Recovering Hyper-V virtual machines to vCloud Director without adding Hyper-V server (*continued*)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

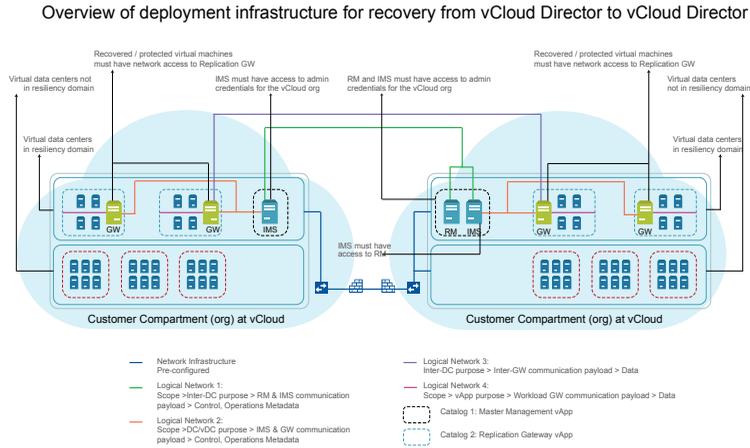
Recovering virtual machines from vCloud Director to vCloud Director

Using Veritas Resiliency Platform , you can configure and protect your virtual machines for recovery from vCloud Director to vCloud Director using the Resiliency Platform Data Mover.

Before starting the product deployment in your data center, ensure that the cloud tenant is created for you and you have the cloud credentials to access it.

Recovering virtual machines from vCloud Director to vCloud Director

Figure 3-11 Overview of deployment infrastructure for recovery from vCloud Director to vCloud Director



The following table provides the summary for deployment, configuration, and recovery of virtual machines from a vCloud Director data center to a vCloud Director data center . These operations can be performed by the end user or by the service subscriber.

Table 3-11 Recovering virtual machines from vCloud Director to vCloud Director

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 3-11 Recovering virtual machines from vCloud Director to vCloud Director *(continued)*

Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances.</p> <p>Download and deploy the virtual appliances on source as well as on the target cloud data center.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ Deploy the virtual appliances for Infrastructure Management Server (IMS) and Replication Gateway in vCloud Director on both the cloud data centers. Resiliency Manager should be deployed either on source or on target data center. If you have multiple virtual data centers, deploy Resiliency Manager , IMS and Replication Gateway in one virtual data center and only IMS and Replication Gateway in rest of the virtual data centers: <ul style="list-style-type: none"> ■ About deploying the virtual appliances ■ Using vCloud Director ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS ■ Configuring Replication Gateways
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add another cloud data center ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Prepare host for replication

Table 3-11 Recovering virtual machines from vCloud Director to vCloud Director *(continued)*

Tasks	More information
Infrastructure Pairing	<p>For recovering assets from vCloud Director to vCloud Director you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering from vCloud Director to vCloud Director.
	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity.</p> <p>You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Manage resiliency groups for remote recovery
	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
	<p>Once you have organized your assets into resiliency groups, you can perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Migrate ■ Take over ■ Resync <p>Note that, Rehearsal and Cleanup Rehearsal operations are not supported for recovery from vCloud Director to vCloud Director.</p>

Table 3-11 Recovering virtual machines from vCloud Director to vCloud Director *(continued)*

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

Recovering VMware virtual machines to Orange Recovery Engine

Using Veritas Resiliency Platform, you can recover VMware virtual machines to Orange Recovery Engine.

The following table provides the summary for deployment, configuration, and recovery of virtual machines to a data center on Orange Recovery Engine.

Table 3-12 Recovering VMware virtual machines to Orange Recovery Engine

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the <i>Veritas Resiliency Platform Overview and Planning Guide</i> to know about the product, its components, features, and capabilities. Refer to the <i>Veritas Resiliency Platform Release Notes</i> for release information such as main features, known issues, and limitations.</p> <p>Overview and Planning Guide</p> <p>Release Notes</p> <p>Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Checklist for deployment and disaster recovery configuration
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the Orange Recovery Engine data center as well as in the premises data center.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway in the Orange Recovery Engine data center: <ul style="list-style-type: none"> ■ Using Orange Recovery Engine ■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> ■ See “Deploying the virtual appliance through VMware vSphere Client” on page 419. ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ See “About configuring the Resiliency Platform components” on page 437. ■ See “Configuring the Resiliency Manager or IMS” on page 439. ■ See “Configuring the Replication Gateways” on page 445.
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings

Table 3-12 Recovering VMware virtual machines to Orange Recovery Engine
(continued)

Tasks	More information
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add VMware servers ■ Prepare host for replication
<p>Infrastructure Pairing</p>	<p>For recovering assets to Orange Recovery Engine you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects ■ For DNS customization, refer Add DNS servers ■ For NIC teaming / bonding, refer Support for NIC teaming / bonding for physical machines ■ Create network mappings, refer Network pairs for recovering virtual machines to Orange Recovery Engine
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for remote recovery.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Configuring resiliency groups for recovery to Orange Recovery Engine
<p>Advance features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual Business Services ■ Resiliency Plans ■ Evacuation Plans

Table 3-12 Recovering VMware virtual machines to Orange Recovery Engine
(continued)

Tasks	More information
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ▪ Rehearsal ▪ Cleanup Rehearsal ▪ Migrate ▪ Recover ▪ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risk ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ▪ Using Klish ▪ Troubleshooting ▪ Updating ▪ References

Recovering physical machines to AWS using Resiliency Platform Data Mover

Using Veritas Resiliency Platform, you can recover physical machines to AWS using Resiliency Platform Data Mover.

The following table provides the summary for deployment, configuration, and recovery of physical machines to a data center on AWS.

Table 3-13 Recovering physical machines to AWS using Resiliency Platform Data Mover

Tasks	More information
<p data-bbox="126 354 360 378">Plan your environment</p> 	<p data-bbox="415 354 1216 465">Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> <li data-bbox="415 486 727 510">■ Overview and Planning Guide <li data-bbox="415 520 585 545">■ Release Notes <li data-bbox="415 555 1018 579">■ Checklist for deployment and disaster recovery configuration
<p data-bbox="126 631 384 687">Deploy and configure the virtual appliances</p> 	<p data-bbox="415 631 1200 715">Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the AWS cloud data center as well as in the premises data center.</p> <ul style="list-style-type: none"> <li data-bbox="415 736 852 760">■ Download the files required for deployment <li data-bbox="415 770 805 795">■ About deploying the virtual appliances <li data-bbox="415 805 1210 854">■ Deploy the Resiliency Platform components in AWS by using one of the following methods: <ul style="list-style-type: none"> <li data-bbox="444 864 1045 888">■ Through AWS marketplace using CloudFormation templates <li data-bbox="444 899 626 923">■ Using OVA files <li data-bbox="415 933 1216 982">■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> <li data-bbox="444 992 760 1017">■ Using VMware vSphere client <li data-bbox="415 1027 1205 1076">■ Deploy Data Gateway in AWS environment if you want to use Object Storage for replication: <ul style="list-style-type: none"> <li data-bbox="444 1086 686 1111">■ Deploy Data Gateway <li data-bbox="415 1121 1161 1145">■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> <li data-bbox="444 1156 852 1180">■ About configuring the virtual appliances <li data-bbox="444 1190 603 1215">■ Prerequisites <li data-bbox="444 1225 848 1249">■ Configuring Resiliency Manager or IMS <li data-bbox="444 1260 801 1284">■ Configuring Replication Gateways

Table 3-13 Recovering physical machines to AWS using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Add Data Gateway (only if you want to use Object Storage mode of replication) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Prepare host for replication
<p>Infrastructure Pairing</p>	<p>For recovering assets to AWS you have to do following infrastructure pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ Create Network group of Cloud Subnets, refer Add network groups (Optional). ■ For DNS customization, refer Add DNS servers. ■ For NIC teaming / bonding, refer Support for NIC teaming / bonding for physical machines ■ For creating network mapping, refer Network pairs for recovering virtual machines to AWS

Table 3-13 Recovering physical machines to AWS using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity.</p> <ul style="list-style-type: none"> ■ Managing physical machines for remote recovery (DR) using Resiliency Platform Data Mover
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities

Table 3-13 Recovering physical machines to AWS using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p data-bbox="126 354 388 378">Miscellaneous references</p> 	<p data-bbox="412 354 1214 466">After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.</p> <ul style="list-style-type: none"> <li data-bbox="412 486 552 510">▪ Using klish <li data-bbox="412 520 599 545">▪ Troubleshooting <li data-bbox="412 555 532 579">▪ Updating <li data-bbox="412 590 556 614">▪ References

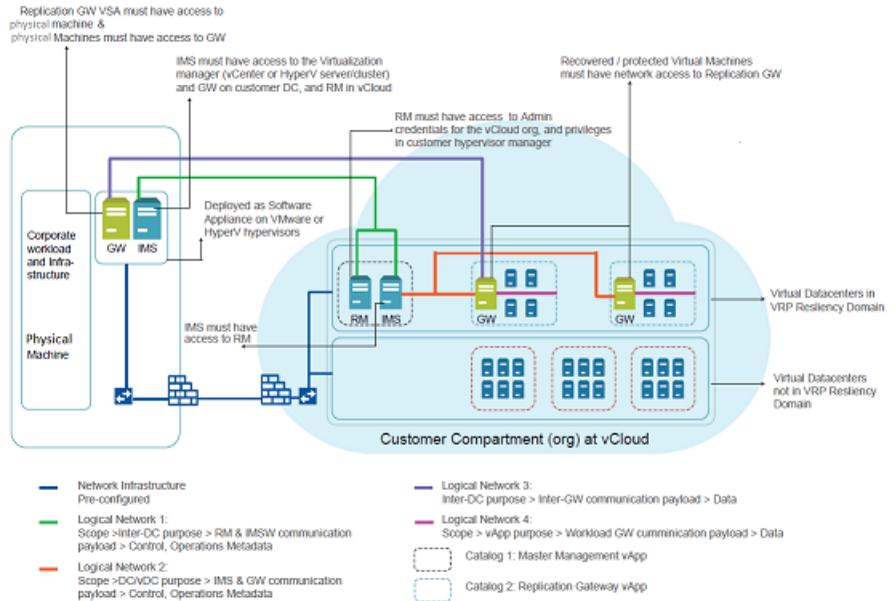
Recovering physical machines to vCloud Director using Resiliency Platform Data Mover

Using Veritas Resiliency Platform, you can you can recover physical machines to vCloud Director using Resiliency Platform Data Mover.

Before starting the product deployment in your data center, ensure that the cloud tenant is created for you and you have the cloud credentials to access it.

Recovering physical machines to vCloud Director using Resiliency Platform Data Mover

Figure 3-12 Overview of deployment infrastructure for recovery to vCloud Director



The following table provides the summary for deployment, configuration, and recovery of physical machines to a data center on vCloud Director. End user or the service subscriber can perform these operations.

Table 3-14 Recovering machines to vCloud Director using Resiliency Platform Data Mover

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment is compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 3-14 Recovering machines to vCloud Director using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p data-bbox="126 354 387 406">Deploy and configure the virtual appliances</p> 	<p data-bbox="413 354 1201 435">Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the AWS cloud data center as well as in the premises data center.</p> <ul style="list-style-type: none"> <li data-bbox="413 456 854 479">■ Download the files required for deployment <li data-bbox="413 487 807 510">■ About deploying the virtual appliances <li data-bbox="413 522 1217 635">■ Deploy the virtual appliances in vCloud Director for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway. If you have multiple virtual data centers, deploy Resiliency Manager and IMS in one virtual data center and only IMS in rest of the virtual data centers: <ul style="list-style-type: none"> <li data-bbox="444 644 686 666">■ Using vCloud Director <li data-bbox="413 678 1217 730">■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> <li data-bbox="444 739 760 762">■ Using VMware vSphere client <li data-bbox="413 774 1163 796">■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> <li data-bbox="444 805 852 828">■ About configuring the virtual appliances <li data-bbox="444 836 852 859">■ Configuring Resiliency Manager or IMS <li data-bbox="444 868 801 890">■ Configuring Replication Gateways
<p data-bbox="126 923 337 975">Set up the resiliency domain</p> 	<p data-bbox="413 923 1217 1005">Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> <li data-bbox="413 1025 982 1048">■ Create the resiliency domain using getting started wizard <li data-bbox="413 1057 899 1079">■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> <li data-bbox="444 1088 561 1111">■ Add IMS <li data-bbox="444 1119 729 1142">■ Add Replication Gateways <li data-bbox="444 1150 1080 1173">■ Add cloud data center (if not done during getting started wizard) <li data-bbox="444 1182 892 1204">■ Manage user authentication and permission <li data-bbox="444 1215 995 1237">■ Manage alerts, notifications, and other product settings
<p data-bbox="126 1279 374 1302">Add asset infrastructure</p> 	<p data-bbox="413 1279 1217 1361">Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> <li data-bbox="413 1381 704 1404">■ Prepare host for replication

Table 3-14 Recovering machines to vCloud Director using Resiliency Platform Data Mover (*continued*)

Tasks	More information
Infrastructure Pairing	<p>For recovering assets to vCloud Director you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ Create Network group of vLAN/Port Group, refer Add network groups (Optional). ■ For DNS customization, refer Add DNS servers. ■ For NIC teaming / bonding, refer Support for NIC teaming / bonding for physical machines ■ Create network mappings, refer Network pairs for recovering virtual machines to vCloud Director.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity.</p> <ul style="list-style-type: none"> ■ Configure physical machines for recovery to on-premises data center
Advanced features 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
Perform remote recovery operations 	<p>Once you have organized your assets into resiliency groups, you can perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Migrate ■ Recover ■ Resync

Table 3-14 Recovering machines to vCloud Director using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

Recovering physical machines to Orange Recovery Engine using Resiliency Platform Data Mover

Using Veritas Resiliency Platform, you can recover physical machines to Orange Recovery Engine using Resiliency Platform Data Mover.

The following table provides the summary for deployment, configuration, and recovery of physical machines to Orange Recovery Engine using Resiliency Platform Data Mover.

Recovering physical machines to Orange Recovery Engine using Resiliency Platform Data Mover

Table 3-15 Recovering physical machines to Orange Recovery Engine using Resiliency Platform Data Mover

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment is compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the Orange Recovery Engine cloud data center as well as in the premises data center.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the Resiliency Platform components in Orange cloud data center using one of the following methods: <ul style="list-style-type: none"> ■ Using Orange Recovery Engine ■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> ■ Using VMware vSphere client ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Prerequisites ■ Configuring Resiliency Manager or IMS ■ Configuring Replication Gateways
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings

Recovering physical machines to Orange Recovery Engine using Resiliency Platform Data Mover

Table 3-15 Recovering physical machines to Orange Recovery Engine using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p data-bbox="126 354 373 378">Add asset infrastructure</p> 	<p data-bbox="413 354 1213 435">Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul data-bbox="413 456 704 480" style="list-style-type: none"> ■ Prepare host for replication
<p data-bbox="126 628 344 652">Infrastructure Pairing</p>	<p data-bbox="413 628 1063 652">For recovering assets to VMware, do following Infrastructure Pairing:</p> <ul data-bbox="413 673 1213 944" style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ For NIC teaming / bonding , refer Support for NIC teaming / bonding for physical machines ■ Create network mappings, refer Network pairs for recovering physical machines to Orange Recovery Engine.
<p data-bbox="126 972 377 996">Create resiliency groups</p> 	<p data-bbox="413 972 1213 1029">After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity.</p> <ul data-bbox="413 1050 1213 1098" style="list-style-type: none"> ■ Managing physical machines for remote recovery (DR) using Resiliency Platform Data Mover
<p data-bbox="126 1246 319 1270">Advanced features</p> 	<p data-bbox="413 1246 1213 1328">Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul data-bbox="413 1348 682 1442" style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans

Table 3-15 Recovering physical machines to Orange Recovery Engine using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ▪ Rehearsal ▪ Cleanup rehearsal ▪ Migrate ▪ Recover ▪ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

Recovering physical machines to Azure using Resiliency Platform Data Mover

Using Veritas Resiliency Platform, you can recover physical machines to Azure using Resiliency Platform Data Mover.

The following table provides the summary for deployment, configuration, and recovery of physical machines to Azure using Resiliency Platform Data Mover.

Recovering physical machines to Azure using Resiliency Platform Data Mover**Table 3-16** Recovering physical machines to Azure using Resiliency Platform Data Mover

Tasks	More information
<p data-bbox="126 354 360 378">Plan your environment</p> 	<p data-bbox="413 354 1210 465">Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> <li data-bbox="413 486 727 510">■ Overview and Planning Guide <li data-bbox="413 520 585 545">■ Release Notes <li data-bbox="413 555 1018 579">■ Checklist for deployment and disaster recovery configuration
<p data-bbox="126 631 384 687">Deploy and configure the virtual appliances</p> 	<p data-bbox="413 631 1204 715">Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in the Azure cloud data center as well as in the premises data center.</p> <ul style="list-style-type: none"> <li data-bbox="413 736 854 760">■ Download the files required for deployment <li data-bbox="413 770 807 795">■ About deploying the virtual appliances <li data-bbox="413 805 1210 923">■ Deploy the virtual appliances for Resiliency Manager, Infrastructure Management Server (IMS), and Replication Gateway in the Azure cloud data center: <ul style="list-style-type: none"> <li data-bbox="444 864 776 888">■ Deploy using Azure PowerShell <li data-bbox="444 899 743 923">■ Through Azure Marketplace <li data-bbox="413 933 1210 1017">■ Deploy the virtual appliances for one or more IMS and Replication Gateway in the premises data center: <ul style="list-style-type: none"> <li data-bbox="444 989 760 1013">■ Using VMware vSphere client <li data-bbox="413 1027 1163 1183">■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> <li data-bbox="444 1055 852 1079">■ About configuring the virtual appliances <li data-bbox="444 1090 602 1114">■ Prerequisites <li data-bbox="444 1124 852 1149">■ Configuring Resiliency Manager or IMS <li data-bbox="444 1159 801 1183">■ Configuring Replication Gateways

Table 3-16 Recovering physical machines to Azure using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Add cloud data center (if not done during getting started wizard) ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings <p>Using the Resiliency Platform console, you can add one or more Azure Stack private cloud instances to the non-cloud datacenter (premise) at source, or at target or both data centers. Azure Stack private cloud can be added to non-cloud datacenter only.</p> <p>Adding Azure Stack private cloud configuration</p>
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Prepare host for replication
<p>Infrastructure Pairing</p>	<p>For recovering assets to Azure you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering physical machines to Azure

Table 3-16 Recovering physical machines to Azure using Resiliency Platform Data Mover (*continued*)

Tasks	More information
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity.</p> <ul style="list-style-type: none"> ■ Managing physical machines for remote recovery (DR) using Resiliency Platform Data Mover
<p>Advance features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities

Table 3-16 Recovering physical machines to Azure using Resiliency Platform Data Mover (*continued*)

Tasks	More information
Miscellaneous references 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update product_name_short components.</p> <ul style="list-style-type: none">■ Using klish■ Troubleshooting■ Updating■ References

Recovery to on-premises data center

This chapter includes the following topics:

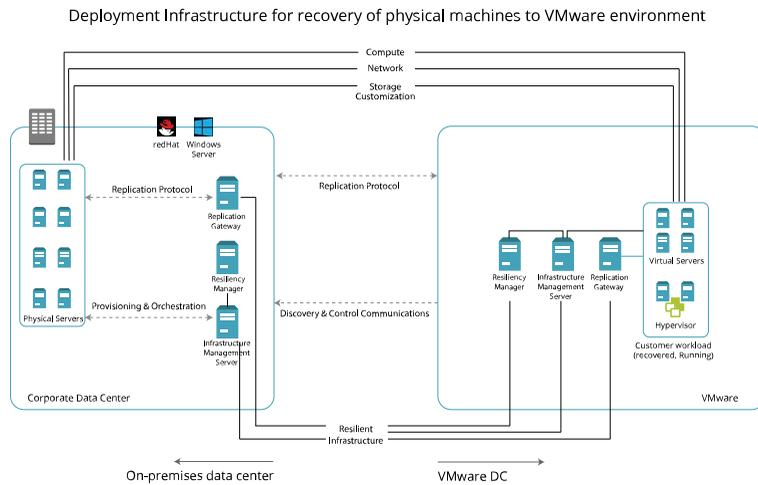
- [Recovering physical machines to VMware virtual machines on an on-premises data center using Resiliency Platform Data Mover](#)
- [Recovering VMware virtual machines to on-premises data center using Resiliency Platform Data Mover](#)
- [Recovering VMware virtual machines from VMware to VMware using NetBackup](#)
- [Recovering virtual machines from VMware to AWS using NetBackup Image Sharing](#)
- [Recovering VMware virtual machines using third-party replication technology](#)
- [Recovering Hyper-V virtual machines using third-party replication technology](#)
- [Recovering Applications using third-party replication technology](#)
- [Recovering InfoScale applications](#)

Recovering physical machines to VMware virtual machines on an on-premises data center using Resiliency Platform Data Mover

Using Veritas Resiliency Platform, you can recover physical machines to VMware virtual machines on an on-premises data center using Resiliency Platform Data Mover.

Note: SD card and USB disks on physical hosts with Veritas Resiliency Platform data mover are not supported.

Figure 4-1 Overview of deployment Infrastructure for recovery of physical machines to VMware virtual machines



The following table provides the summary for deployment, configuration, and recovery of physical machines to on-premises data center using Resiliency Platform Data Mover.

Table 4-1 Recovering physical machines DC on on-premises data center using Resiliency Platform Data Mover

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 4-1 Recovering physical machines to on-premises data center using Resiliency Platform Data Mover (*continued*)

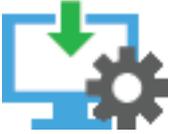
Tasks	More information
<p data-bbox="126 354 385 409">Deploy and configure the virtual appliances</p> 	<p data-bbox="412 354 1210 435">Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances for Resiliency Manager, IMS, and Replication Gateway in both the data centers.</p> <ul style="list-style-type: none"> <li data-bbox="412 456 854 479">■ Download the files required for deployment <li data-bbox="412 487 807 510">■ About deploying the virtual appliances <li data-bbox="412 520 1002 543">■ Deploy the virtual appliances using VMware vSphere client <li data-bbox="412 553 1161 576">■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> <li data-bbox="444 586 852 609">■ About configuring the virtual appliances <li data-bbox="444 619 850 642">■ Configuring Resiliency Manager or IMS <li data-bbox="444 652 801 675">■ Configuring Replication Gateways
<p data-bbox="126 706 337 762">Set up the resiliency domain</p> 	<p data-bbox="412 706 1217 788">Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> <li data-bbox="412 808 982 831">■ Create the resiliency domain using getting started wizard <li data-bbox="412 841 897 864">■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> <li data-bbox="444 874 561 897">■ Add IMS <li data-bbox="444 907 729 930">■ Add Replication Gateways <li data-bbox="444 940 1174 963">■ Configuring Replication Gateway as a PXE Boot server and DHCP server <li data-bbox="444 973 892 996">■ Manage user authentication and permission <li data-bbox="444 1006 995 1029">■ Manage alerts, notifications, and other product settings
<p data-bbox="126 1058 374 1081">Add asset infrastructure</p> 	<p data-bbox="412 1058 1217 1140">Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> <li data-bbox="412 1161 770 1183">■ Add VMware virtualization servers <li data-bbox="412 1194 704 1216">■ Prepare host for replication

Table 4-1 Recovering physical machines to on-premises data center using Resiliency Platform Data Mover (*continued*)

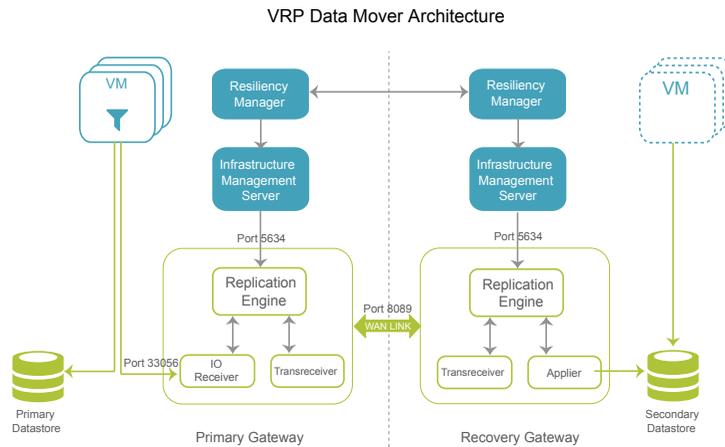
Tasks	More information
Infrastructure Pairing	<p>For recovering assets to VMware you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ For NIC teaming / bonding, refer Support for NIC teaming / bonding for physical machines ■ Create network mappings, refer Network pairs for recovering physical machines to VMware.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity.</p> <ul style="list-style-type: none"> ■ Configure physical machines for recovery to on-premises data center
Monitor assets 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities
Miscellaneous references 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.</p> <ul style="list-style-type: none"> ■ Using klish ■ Troubleshooting ■ Updating ■ References

Recovering VMware virtual machines to on-premises data center using Resiliency Platform Data Mover

Using Veritas Resiliency Platform, you can recover VMware virtual machine to on-premises data center using Resiliency Platform Data Mover. For recovering VMware virtual machines to on-premises data center, Resiliency Platform Data Mover uses VMware VAIO (vSphere APIs for IO Filter) interfaces published and supported by VMware.

Veritas Resiliency Platform now supports Continuous Data Protection (CDP) mechanism to recover from the current or a point in past known as Recovery Point. This support is available from version 3.5 and is now supported only for recovering VMware virtual machines to on-premises data center using Resiliency Platform Data Mover using VMware VAIO (vSphere APIs for IO Filter) interfaces. CDP is enabled when you create or edit a resiliency group while selecting the Replication Gateway pair.

Figure 4-2 Overview of deployment Infrastructure for recovery using Resiliency Platform Data Mover



The following table provides the summary for deployment, configuration, and recovery of VMware virtual machines to on-premises data center using data mover.

Table 4-2 Recovering VMware virtual machines using VMware VAIO

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances for Resiliency Manager, IMS, and Replication Gateway in both the data centers.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances using VMware vSphere client ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS ■ Configuring Replication Gateways
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add Replication Gateways ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings ■ For secure communication with VMware vCenter server, install the root CA certificate. Refer Managing security
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add VMware virtualization servers

Table 4-2 Recovering VMware virtual machines using VMware VAIO
(continued)

Tasks	More information
Infrastructure Pairing	<p>For recovering assets to VMware you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Infrastructure Pairing > Replication Appliance, refer Create Replication Gateway pair. ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering machines to on-premises data center.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for recovery to remote data center.</p> <p>From version 3.5, CDP storage mechanism is provided to the resiliency groups. While configuring the resiliency group for disaster recovery, in the Replication Gateway pair selection wizard, you can enable the CDP storage and can provide the % of the CDP storage to be used on the source and target data centers.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for monitoring ■ Configure VMware virtual machines for recovery to on-premises data center
Advanced features 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
Perform remote recovery operations 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal See "Performing cleanup rehearsal for virtual machines" on page 535. ■ Migrate ■ Recover ■ Resync

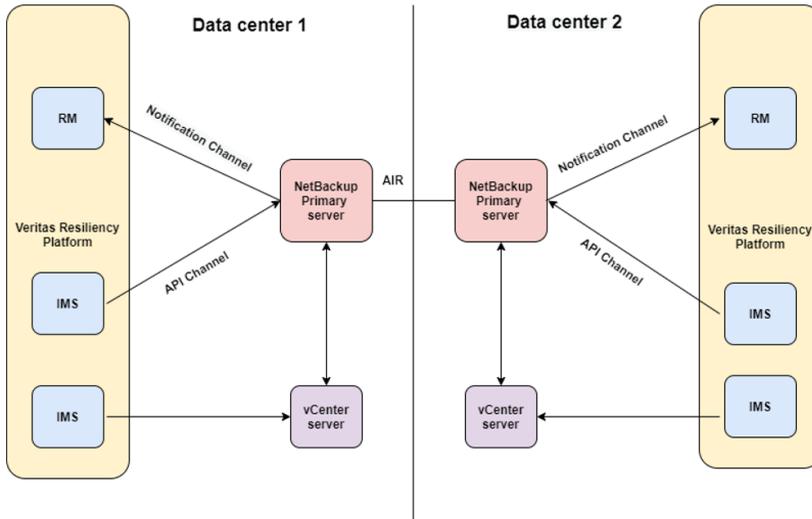
Table 4-2 Recovering VMware virtual machines using VMware VAIO
(continued)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ▪ Risks ▪ Reports ▪ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. You can also use klish to update Resiliency Platform components.</p> <ul style="list-style-type: none"> ▪ Using klish ▪ Troubleshooting ▪ Updating ▪ References

Recovering VMware virtual machines from VMware to VMware using NetBackup

Using the Veritas Resiliency Platform 10.0, you can restore VMware virtual machine from NetBackup generated backup images to the target data center. For more information on NetBackup and NetBackup Appliances, see [About NetBackup and NetBackup Appliances](#).

Figure 4-3 Deployment architecture for NetBackup primary server



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

The following table provides the summary for deployment, configuration, and recovery of virtual machines from NetBackup generated backup images.

Table 4-3 Recovering virtual machines using NetBackup images

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the <i>Veritas Resiliency Platform Overview and Planning Guide</i> to know about the product, its components, features, and capabilities. Refer to the <i>Veritas Resiliency Platform Release Notes</i> for release information such as main features, known issues, and limitations.</p> <p>Overview and Planning Guide</p> <p>Release Notes</p> <p>Ensure that the configuration details in your environment matches the requirements mentioned in the checklist.</p>

Table 4-3 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances for Resiliency Manager and IMS in both the data centers.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances using VMware vSphere client ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings ■ For secure communication, refer Managing security
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add VMware servers ■ Add NetBackup primary server ■ Add IMS to NetBackup primary server as an additional server
<p>Infrastructure Pairing</p>	<p>For recovering assets to VMware you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering machines to on-premises data center.

Table 4-3 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for recovery on local or remote data center.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Manage VMware virtual machines for remote recovery using NetBackup images
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
<p>Perform recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform recover (local or remote) operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Recover virtual machines
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities

Table 4-3 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ■ Using klish ■ Troubleshooting ■ Updating ■ References

Recovering virtual machines from VMware to AWS using NetBackup Image Sharing

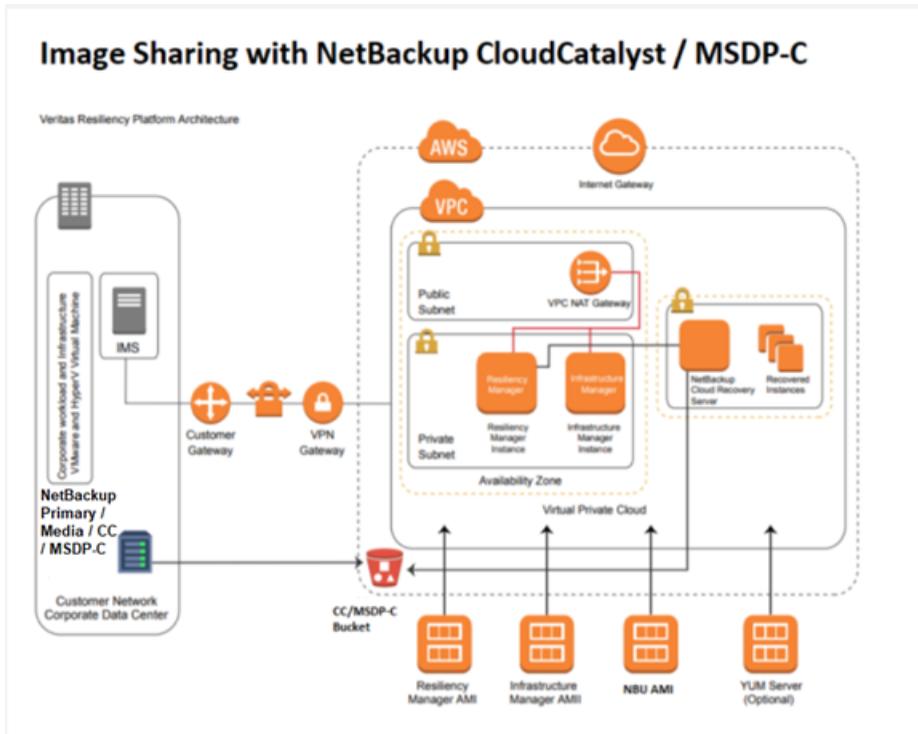
Using the Resiliency Platform, you can recover VMware virtual machine from NetBackup generated backup images that are stored into AWS S3 buckets to the AWS cloud target data center.

In figure the on-premises data center is the source data center and the target data center is an AWS cloud region. The Infrastructure Management Server (IMS) in the on-premises data center discovers the vCenter server and the backup configuration from the NetBackup primary server. NetBackup with the Image Sharing feature, which is available from NetBackup version 8.2 onwards, backs up the virtual machine images along with the image metadata from the on-premise data center into the designated S3 bucket.

For backing up the data in AWS S3 bucket, either NetBackup CloudCatalyst or MSDP-C configuration can be used.

The recovery using these backup images is achieved using a NetBackup Cloud Recovery Server (CRS) virtual appliance that is deployed in the cloud data center. The image metadata stored in the S3 bucket allows the NetBackup CRS to read the image information and create an Amazon Machine Image (AMI). This AMI is then used to provision cloud instances in the cloud data center during the recover operation.

Figure 4-4 Image Sharing with NetBackup CloudCatalyst / MSDP-C



The following table provides the summary of deployment, configuration, and recovery of virtual machines to cloud using NetBackup generated backup images that are stored in S3 bucket using the Image Sharing feature.

Table 4-4 Recovering virtual machines using NetBackup images

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the <i>Veritas Resiliency Platform Overview and Planning Guide</i> to know about the product, its components, features, and capabilities. Refer to the <i>Veritas Resiliency Platform Release Notes</i> for release information such as main features, known issues, and limitations.</p> <p>Overview and Planning Guide</p> <p>Release Notes</p> <p>Ensure that the configuration details in your environment match the requirements mentioned in the checklist.</p> <p>Checklist for recovery of VMware virtual machines to AWS cloud using NetBackup</p>

Table 4-4 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances for Resiliency Manager and IMS in both the data centers.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment Downloading the Veritas Resiliency Platform virtual appliances ■ About deploying the virtual appliances About deploying the Resiliency Platform virtual appliances ■ Deploy the Resiliency Platform components in AWS by using one of the following methods: <ul style="list-style-type: none"> ■ Deploying the virtual appliances in AWS through AWS Marketplace ■ Deploying the virtual appliances in AWS using OVA files ■ Deploy the virtual appliances for one or more IMS in the premises data center: <ul style="list-style-type: none"> ■ Deploying the virtual appliance through VMware vSphere Client ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ See “About configuring the Resiliency Platform components ” on page 437. About configuring the Resiliency Platform components ■ See “Prerequisites for configuring Resiliency Platform components” on page 438. Prerequisites for configuring Resiliency Platform components ■ Deploy the Cloud Recovery Server (CRS) Deploy CRS <p>For more information on CRS, refer to <i>About Image Sharing using MSDP cloud</i> topic in <i>NetBackup Deduplication Guide</i>.</p>
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Getting started with a new Resiliency Platform configuration ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Adding an IMS ■ Adding AWS cloud data center ■ Managing user authentication and permissions ■ For secure communication, refer Managing security
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Adding VMware virtualization servers ■ Adding NetBackup primary server ■ Adding NetBackup Cloud Recovery Server (CRS)

Table 4-4 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Infrastructure Pairing</p>	<p>For recovering assets to VMware you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ Create Network group of cloud subnets. Adding a network group ■ Customize DNS Configuring DNS server settings for a data center ■ Create network mappings. Network pairs for recovering virtual machines to AWS
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for recovery on local or remote data center.</p> <ul style="list-style-type: none"> ■ Configuring a resiliency group for basic monitoring ■ Managing VMware virtual machines for remote recovery to AWS cloud using NetBackup images
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Managing virtual business services ■ Managing resiliency plans ■ About evacuation plan
<p>Perform recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform recover (local or remote) operations on the resiliency groups.</p> <p>Note: The rehearsal operation is not supported from AWS to VMware.</p> <ul style="list-style-type: none"> ■ Performing the rehearsal operation on virtual machines from VMware to AWS using NetBackup Image Sharing ■ Performing cleanup rehearsal for virtual machines ■ Recovering virtual machines to cloud (AWS) using NetBackup Image Sharing

Table 4-4 Recovering virtual machines using NetBackup images (*continued*)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ About risks ■ About reports ■ Managing activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ■ About klish ■ Troubleshoot ■ About applying updates to Resiliency Platform ■ References

Recovering VMware virtual machines using third-party replication technology

When you configure VMware virtual machines for disaster recovery, Veritas Resiliency Platform lets you select the replication technology to replicate data from a production data center to a recovery data center.

Veritas Resiliency Platform supports the following replication technologies. Depending on your environment, select the replication technology that best fits your business needs.

- EMC SRDF
- EMC Recoverpoint
- Netapp (cDOT) Snapmirror
- HP 3PAR Remote Copy
- Hitachi TrueCopy/HUR
- IBM SVC Global Mirror
- IBM XIV Remote Mirror

Table 4-5 Recovering VMware virtual machines using third-party replication technology

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances for Resiliency Manager and IMS in both the data centers.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances using VMware vSphere client ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add VMware virtualization servers ■ Add enclosures

Table 4-5 Recovering VMware virtual machines using third-party replication technology (*continued*)

Tasks	More information
Infrastructure Pairing	<p>For recovering assets to VMware you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering machines to on-premises data center.
Create resiliency groups 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for recovery on local or remote data center.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Manage resiliency groups for remote recovery
Advanced features 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
Perform remote recovery operations 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync

Table 4-5 Recovering VMware virtual machines using third-party replication technology (*continued*)

Tasks	More information
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ■ Using klish ■ Troubleshooting ■ Updating ■ References

Recovering Hyper-V virtual machines using third-party replication technology

When you configure Hyper-V virtual machines for disaster recovery, Veritas Resiliency Platform lets you select the replication technology to replicate data from a production data center to a recovery data center.

Veritas Resiliency Platform supports the following replication technologies. Depending on your environment, select the replication technology that best fits your business needs.

- Hyper-V Replica
- EMC SRDF
- EMC Recoverpoint
- Netapp (cDOT) Snapmirror
- HP 3PAR Remote Copy
- Hitachi TrueCopy/HUR
- IBM SVC Global Mirror

- IBM XIV Remote Mirror
- Infinidat

Table 4-6 Recovering Hyper-V virtual machines using third-party replication technology

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in both the data centers.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances for Resiliency Manager and Infrastructure Management Server (IMS) <ul style="list-style-type: none"> ■ Using VMware vSphere client ■ Using Hyper-V Manager ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings

Table 4-6 Recovering Hyper-V virtual machines using third-party replication technology (*continued*)

Tasks	More information
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add Hyper-V servers ■ Add enclosures
<p>Infrastructure Pairing</p>	<p>For recovering assets to Hyper-V you have to do following Infrastructure Pairing:</p> <ul style="list-style-type: none"> ■ Navigate to Settings > Infrastructure > Access Profile > Network to mark purpose of the networks, refer Add and map network objects. ■ For DNS customization, refer Add DNS servers. ■ Create network mappings, refer Network pairs for recovering machines to on-premises data center.
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for recovery on local or remote data center.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Manage resiliency groups for remote recovery
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans

Table 4-6 Recovering Hyper-V virtual machines using third-party replication technology (*continued*)

Tasks	More information
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ■ Using klish ■ Troubleshooting ■ Updating ■ References

Recovering Applications using third-party replication technology

Veritas Resiliency Platform supports following replication technology for recovery of the applications:

- DataGuard

Table 4-7 Recovering applications using third-party replication technology

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances in both the data centers.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances for Resiliency Manager and Infrastructure Management Server (IMS) <ul style="list-style-type: none"> ■ Using VMware vSphere client ■ Using Hyper-V Manager ■ Configure the virtual appliances as Veritas Resiliency Platform components: <ul style="list-style-type: none"> ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Add asset infrastructure</p> 	<p>Before you can monitor and manage data center assets from the console, you must add the asset infrastructure to Veritas Resiliency Platform. The IMS then discovers the asset information for monitoring and operations in the console.</p> <ul style="list-style-type: none"> ■ Add virtualization servers: <ul style="list-style-type: none"> ■ Add VMware virtualization servers ■ Hyper-V servers ■ Add host assets ■ Add enclosures ■ Add DNS servers

Table 4-7 Recovering applications using third-party replication technology
(continued)

Tasks	More information
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for recovery on local or remote data center.</p> <ul style="list-style-type: none"> ■ Managing applications ■ Configure resiliency groups for basic monitoring ■ Manage applications for remote recovery
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities

Table 4-7 Recovering applications using third-party replication technology
(continued)

Tasks	More information
Miscellaneous references 	After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot. <ul style="list-style-type: none">▪ Using klish▪ Troubleshooting▪ Updating▪ References

Recovering InfoScale applications

Recovering InfoScale applications

Veritas Resiliency Platform lets you manage the InfoScale applications by configuring the corresponding clusters into Infrastructure Management Servers (IMS). The InfoScale applications are automatically discovered in the Resiliency Platform. You can group the InfoScale applications into resiliency groups or VBS to recover, monitor, visualize, and generate reports about these applications in the Resiliency Platform. Before version 3.5, you were required to add Veritas InfoScale Operations Manager into Veritas Resiliency Platform to manage the InfoScale applications. From version 3.5 there is no need to use Veritas InfoScale Operations Manager to manage the InfoScale applications. If you are upgrading from version 3.4 or earlier version where you had configured Veritas InfoScale Operations Manager, then refer to topic for how to configure those clusters into Infrastructure Management Servers.

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

Figure 4-5 A typical workflow for recovering managed InfoScale applications

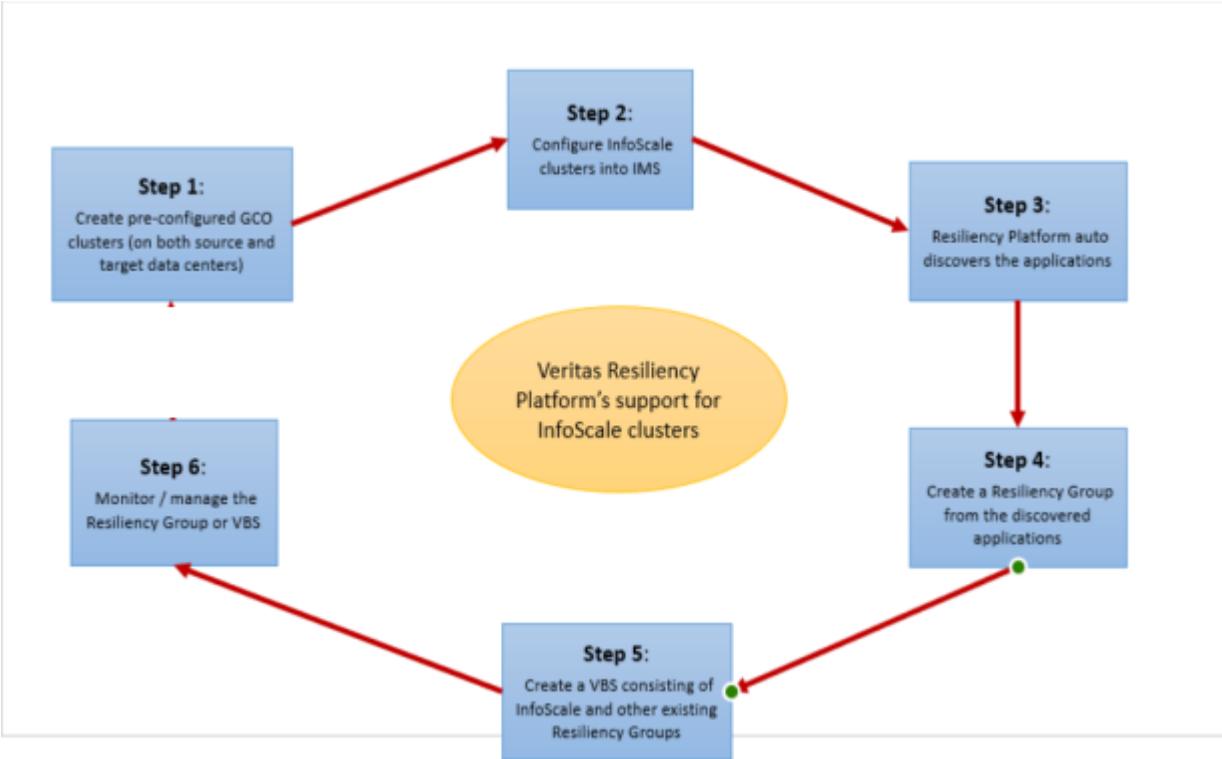


Table 4-8 Recovering InfoScale applications

Tasks	More information
<p>Plan your environment</p> 	<p>Refer to the Overview and Planning Guide to know about the product, its components, features, and capabilities. Refer to the Release Notes for release information such as main features, known issues, and limitations. Ensure that the configuration details in your environment are compatible with the requirements mentioned in the checklist.</p> <ul style="list-style-type: none"> ■ Overview and Planning Guide ■ Release Notes ■ Checklist for deployment and disaster recovery configuration

Table 4-8 Recovering InfoScale applications (*continued*)

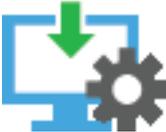
Tasks	More information
<p>Deploy and configure the virtual appliances</p> 	<p>Veritas Resiliency Platform is deployed as virtual appliances. Download and deploy the virtual appliances for Resiliency Manager and IMS in both the data centers.</p> <ul style="list-style-type: none"> ■ Download the files required for deployment ■ About deploying the virtual appliances ■ Deploy the virtual appliances for Resiliency Manager and Infrastructure Management Server (IMS) <ul style="list-style-type: none"> ■ Using VMware vSphere client ■ Using Hyper-V Manager ■ About configuring the virtual appliances ■ Configuring Resiliency Manager or IMS
<p>Set up the resiliency domain</p> 	<p>Set up the infrastructure and basic settings of the Veritas Resiliency Platform resiliency domain. These tasks are performed after you configure the Resiliency Manager and log in to the web console.</p> <ul style="list-style-type: none"> ■ Create the resiliency domain using getting started wizard ■ Configure the settings for the resiliency domain: <ul style="list-style-type: none"> ■ Add IMS ■ Add InfoScale cluster ■ Manage user authentication and permission ■ Manage alerts, notifications, and other product settings
<p>Create resiliency groups</p> 	<p>After adding the assets to Resiliency Platform, you organize the related assets into a resiliency group that you can protect and manage as a single entity. You can create a resiliency group either for basic monitoring (start or stop virtual machines) or for recovery on local or remote data center.</p> <ul style="list-style-type: none"> ■ Configure resiliency groups for basic monitoring ■ Manage applications for remote recovery
<p>Advanced features</p> 	<p>Virtual business services, resiliency plans, and evacuation plans are some of the features of Veritas Resiliency Platform that you can additionally use to customize the process of recovery of your assets.</p> <ul style="list-style-type: none"> ■ Virtual business services ■ Resiliency plans ■ Evacuation plans

Table 4-8 Recovering InfoScale applications (*continued*)

Tasks	More information
<p>Perform remote recovery operations</p> 	<p>Once you have configured the resiliency groups for remote recovery, you can perform rehearsal and cleanup rehearsal operations on the resiliency groups. You can also perform migrate, recover, or resync operations on the resiliency groups.</p> <ul style="list-style-type: none"> ■ Rehearsal ■ Cleanup rehearsal ■ Migrate ■ Recover ■ Resync
<p>Monitor assets</p> 	<p>You can monitor risks to the recoverability or continuity of your protected assets. Run various reports to view the status of the assets in your data center. And view details about operations such as the status (in-progress, finished, or failed), start and end time, and the objects on which the operation was performed on the Activities page.</p> <ul style="list-style-type: none"> ■ Risks ■ Reports ■ Activities
<p>Miscellaneous references</p> 	<p>After the virtual appliances are deployed and configured, you are given limited menu-based access to the operating system and the product. You need to use klish menu to manage the configuration-related changes to the product and to troubleshoot.</p> <ul style="list-style-type: none"> ■ Using klish ■ Troubleshooting ■ Updating ■ References

Index

A

- activities
 - abort 270, 582
- alert settings
 - email 156, 484
 - SNMP 158, 486
- asset infrastructure
 - about adding assets as hosts 191, 503
 - adding Hyper-V virtualization servers 495
 - adding VMware virtualization servers 167
 - prerequisites for hosts 193, 505
 - prerequisites for Hyper-V virtualization discovery 495
 - prerequisites for VMware discovery 169
- assets
 - configuring for monitoring 203, 217, 515, 529
- authentication domains
 - configuring 141, 469
 - editing 147, 475
 - unconfiguring 146, 474

C

- cloud
 - configurations 117
 - editing configuration 132
 - refreshing 133
- cloud configuration
 - removing 134
- cloud data center
 - removing 134
- configuring for remote recovery
 - using Google Cloud Platform 205, 517
- customize panel
 - GCP 209, 520

D

- deploying
 - AWS 22, 37, 374, 389
 - Hyper-V Manager 420
 - Orange Recovery Engine 63, 415

- deploying (*continued*)
 - vCloud 61, 413
 - virtual appliance 16, 368
 - VMware vSphere Client 419
- deploying virtual appliances
 - Azure 44, 396
 - Azure Stack 49, 401
- deployment in AWS
 - CloudFormation 22, 374
 - creating AMI 42, 394
 - launching instance 43, 395
 - OVA 37, 389
 - prerequisites 24, 376
 - uploading OVA through CLI 42, 394
 - uploading OVA through web 41, 393
- deployment in Azure
 - prerequisites 47, 51, 399, 403
- deployment in Orange Recovery Engine 63, 415
- deployment in vCloud 61, 413
- deployment in vCloud Director
 - prerequisites 62, 414
- disaster recovery operations
 - cleanup rehearsal 223, 535
 - migrate 536
 - recover 539
 - rehearse 220, 532
 - resync 225, 537
- disaster recovery to GCP
 - prerequisites for Hyper-V 518
 - prerequisites for VMware 206
- domains on login 166, 494
- downloading
 - virtual appliances 366

E

- email settings 156, 484
- evacuation plan
 - about 218, 530

F

- fresh deployment
 - update Resiliency Platform 370
- from vCloud Director to vCloud Director 712

G

- GCP
 - deploy using OVA files 429
- GCP marketplace
 - deploy virtual appliances 69, 421
- global user
 - configuring 154, 482

H

- host
 - preparing for replication 184, 496
- hosts
 - prerequisites for adding 193, 505
- hosts for replication 184, 496
- Hyper-V servers
 - adding 495
 - prerequisites 495

I

- image using command - based method 435
- image using web-based 435
- IMS
 - configuring 92, 439
- Infrastructure Management Server
 - adding 109, 456
- inputs for template 71, 423

J

- jobs for custom personas 151, 479

K

- klish
 - about 272, 584

L

- launch virtual appliances 436
- log files
 - downloading 166, 494
- login 166, 494
- logs
 - purge settings 165, 493

N

- network customization
 - options 211, 523
- network mapping in GCP 200, 512
- network objects
 - about 197, 509
- network pair
 - creating 202, 514
- notification settings
 - email 156, 484
 - rules 162, 490
 - SNMP 158, 486

O

- OVA file
 - command-line method 87, 434

P

- permissions 134, 463
 - assigning to users 148, 476
 - overview 136, 464
- personas
 - custom 150–151, 478–479
 - limiting object scope for operations 154, 482
 - predefined 137, 465
- policy statement
 - Orange Recovery Engine 124
- post-upgrade
 - tasks 324, 636
- pre-upgrade
 - checklist 329, 641
- prerequisites
 - GCP 76–77, 428–429
- purge setting
 - logs and SNMP traps 165, 493

R

- recover applications
 - using third-party replication technology 756
- recover Hyper-V
 - to AWS 675
 - to Azure 687
 - to vCloud Director 700
 - to vCloud Director without adding Hyper-V server 708
 - using third-party replication technology 753
- recover InfoScale applications 759

- recover physical machine
 - to AWS 719
 - to Azure using Resiliency Platform Data Mover 730
 - to on-premises data center using Resiliency Platform Data Mover 735
 - to Orange Recovery Engine using Resiliency Platform Data Mover 727
- recover physical machines
 - to vCloud Director 723
- recover virtual machines
 - Azure cloud to Azure cloud 691
 - to vCloud Director 712
- recover VMware
 - to AWS 671
 - to Azure 683
 - to on-premises data center using Resiliency Platform Data Mover 739
 - to Orange Recovery Engine 716
 - to vCloud Director 696
 - to vCloud Director without adding vCenter server 704
 - using NetBackup images 742
 - using third-party replication technology 750
- rehearsal
 - prerequisites 222, 534
- rehearse operations 220, 223, 532, 535
- replace
 - Replication Gateway 660
- Replication Gateway
 - configuring 98, 445
- reports
 - about 264, 576
 - inventory 264, 576
 - risk assessment 264, 576
 - running 267, 579
 - scheduling 265, 577
 - viewing 268, 580
- resiliency domain
 - joining 108, 455
- resiliency groups
 - guidelines 216, 528
- resiliency groups-virtual machines
 - about 216, 528
- Resiliency Manager
 - configuring 92, 439
 - joining existing resiliency domain 108, 455
- risks
 - about 230, 542

- risks *(continued)*
 - description 232, 544
- rules for event notifications 162, 490

S

- SNMP
 - configuring settings 158, 486
 - MIB file 162, 490
 - purge settings for traps 165, 493
- Syslog
 - managing 163, 491

T

- telemetry collection 165, 493

U

- updates
 - about 319, 631
 - Resiliency Managers 349
- upload ova files 86, 433
- user authentication 134, 136, 463–464
- user permissions
 - overview 136, 464
- users
 - assigning permissions 148, 476
 - configuring 147, 475

V

- VBS
 - configuration states 215, 527
- Veritas Resiliency Platform
 - configuring 90, 437
- virtual appliance
 - deploying 16, 368
 - downloading 366
- virtual business services
 - about 526
 - understanding tiers 214, 526
- virtual machine
 - add to IMS 184, 496
- virtual machines
 - configuring for monitoring 203, 217, 515, 529
- VMware vCenter Server privileges
 - cloud replication 177
 - configuring ESX 170
 - NetBackup integration 179
 - Physical machines to VMware 181
 - third party replication 171

- VMware vCenter Server privileges *(continued)*
 - VAIO 173
- VMware virtualization server
 - adding 167
 - requirements for IMS discovery 169
- volume type
 - AWS, GCP 208, 519

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, virtualization servers, virtual machines, enclosures, and applications.
assets	The virtual machines, physical machines, or applications that have been discovered by the Infrastructure Management Server (IMS) and that can be grouped into resiliency groups.
data center	<p>A location that contains asset infrastructure to be managed by Veritas Resiliency Platform.</p> <p>For the disaster recovery use case, the resiliency domain must contain at least two data centers in different locations, a source data center and target data center. Each data center has a Resiliency Manager and one or more IMSs.</p>
host	In Veritas Resiliency Platform, the term hosts means Application host, Resiliency Platform Data Mover host, Storage discovery host, VMware Discovery host, and Hyper-V host.
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.
klish	Command Line Interface SHell. Provides the command line menu on the virtual appliance for use after the initial bootstrap configuration.
migrate	A planned activity involving graceful shutdown of assets at the source data center and starting them at the target data center. In this process, replication ensures that consistent data is made available at the target 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.
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.
Replication Gateway	The Veritas Resiliency Platform component that performs data replication between the source and the target data center.
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 to be 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 web 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.
Resiliency Platform Data Mover Replication host	To enable replication using Resiliency Platform Data Mover replication technology, you need to add an asset and prepare it for replication. Asset can be a physical machine or a virtual machine.
source data center	The data center that is normally used for business.
take over	An activity initiated by a user when the source data center is down due to a disaster and the assets need to be restored at the target data center to provide business continuity.
target data center	The data center that is used if a disaster scenario occurs.
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 operations are performed on the resiliency groups.
VAIO framework	VMware framework consisting of vSphere APIs for I/O Filtering. This framework enables Veritas Resiliency Platform to run filters on ESXi servers and intercept any I/O requests from a guest operating system to a virtual disk.
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 recovery in case of a disaster in the desired order.
Veritas Replication Set	A virtual machine, which belongs to the resiliency group, is termed as Veritas Replication Set. All the disks attached to this virtual machine, including the boot and data disk, constitute a Veritas Replication Set. The write order fidelity is maintained across all disks in a given replication set.
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.