

Veritas™ Resiliency Platform 2.0: Solutions for VMware



Veritas Resiliency Platform: Solutions for VMware

Last updated: 2016-09-13

Document version: 2.0 Rev 1

Legal Notice

Copyright © 2016 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
500 E Middlefield Road
Mountain View, CA 94043

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Overview of Resiliency Platform	9
	About Veritas Resiliency Platform	9
	About disaster recovery using Resiliency Platform	10
	About Resiliency Platform features and components	11
	About Resiliency Platform capabilities	13
	About managing VMware virtual machines using Resiliency Platform	14
	About permissions for operations in the console	14
Chapter 2	Managing resiliency groups	16
	About resiliency groups	16
	Guidelines for organizing resiliency groups	17
	About service objectives	17
	Managing virtual machines for basic monitoring	18
	Starting a resiliency group	19
	Stopping a resiliency group	20
	Displaying resiliency group information and status	21
	Viewing resiliency group details	24
	Editing a resiliency group	24
	Deleting a resiliency group	25
Chapter 3	Overview of Resiliency Platform Data Mover	27
	About Veritas Resiliency Platform Data Mover	27
	Supported environments for Resiliency Platform Data Mover with VMware VAIO	28
	Resiliency Platform Data Mover deployment process overview	28
	How Resiliency Platform Data Mover works	29
	About Veritas Resiliency Platform Data Mover architecture	30
	About full synchronization with Veritas Resiliency Platform Data Mover	32
	How Veritas Resiliency Platform Data Mover handles virtual machine writes	33
	How virtual appliances for Veritas Resiliency Platform Data Mover are deployed and configured	34

Chapter 4	Preparing for disaster recovery configuration	36
	Supported replication technologies with Veritas Resiliency Platform	36
	About the Veritas Resiliency Platform Data Mover Replication Gateway	37
	DR for VMware virtual machines using Resiliency Platform Data Mover - an overview of key steps	37
	DR for VMware virtual machines using array-based replication - an overview of key steps	39
	Configuring VMware virtual machines disaster recovery using EMC SRDF replication	41
	Configuring VMware disaster recovery using NetApp SnapMirror	43
	Configuring VMware virtual machines disaster recovery using EMC RecoverPoint replication	45
	Configuring VMware virtual machines disaster recovery using Hitachi True Copy replication	47
	Configuring VMware virtual machines disaster recovery using HPE 3PAR Remote Copy replication	48
	Prerequisites for configuring VMware virtual machines for disaster recovery	50
	VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover	52
	Limitations for virtual machine disaster recovery	54
Chapter 5	Managing disaster recovery network mapping	55
	Viewing and configuring network settings for a data center	55
	Editing network settings for a data center	56
	Removing network settings for a data center	56
	Configuring DNS server settings for a data center	57
	Setting up network mapping between production and recovery data centers	58
Chapter 6	Managing Resiliency Platform Data Mover gateway pairing	60
	Viewing Veritas Replication Gateways	60
	About Veritas Replication Gateway pairs	61
	How Resiliency Platform Data Mover supports encryption for data replication	61
	Creating a Veritas Replication Gateway pair	61

	Removing a Veritas Replication Gateway pair	62
	Viewing Veritas Replication Gateway pairs	63
	Modifying encryption for a Veritas Replication Gateway pair	63
Chapter 7	Configuring resiliency groups for remote recovery	65
	Understanding the role of resiliency groups in disaster recovery operations	65
	How Resiliency Platform configures disaster recovery protection for virtual machines	66
	Managing virtual machines for remote recovery (DR)	67
	DR configuration options for VMware virtual machines using Resiliency Platform Data Mover	68
	Network customization options	70
	Verifying the replication status for Veritas Resiliency Platform Data Mover	71
	Deleting a resiliency group configured with Resiliency Platform Data Mover	72
Chapter 8	Rehearsing DR operations to ensure DR readiness	74
	About ensuring the disaster recovery readiness of your assets	74
	Rehearse operations - array-based replication	75
	Rehearse operations - Resiliency Platform Data Mover	76
	Performing the rehearsal operation	76
	Performing cleanup rehearsal	77
Chapter 9	Performing disaster recovery operations	79
	How Resiliency Platform Data Mover handles migrate	79
	Migrating a resiliency group of virtual machines	80
	How Resiliency Platform Data Mover handles takeover	81
	How Resiliency Platform Data Mover migrates workloads back to the source data center after a takeover	81
	Taking over a resiliency group of virtual machines	82
	Performing the resync operation	82
Chapter 10	Monitoring and reporting on assets status	84
	About the Resiliency Platform Dashboard	84
	Understanding asset types	86
	Displaying an overview of your assets	86

	Viewing reports	87
Chapter 11	Monitoring risks	89
	About risk insight	89
	Displaying risk information	90
	Predefined risks in Resiliency Platform	91
	Viewing the current risk report	96
	Viewing the historical risk report	97
Chapter 12	Managing activities and resiliency plans	98
	Managing activities	98
	Viewing activities	98
	Aborting a running activity	99
	Managing resiliency plans	100
	About resiliency plans	100
	Creating a new resiliency plan template	101
	Editing a resiliency plan template	104
	Deleting a resiliency plan template	105
	Viewing a resiliency plan template	105
	Creating a new resiliency plan	106
	Editing a resiliency plan	107
	Deleting a resiliency plan	108
	Executing a resiliency plan	108
	Viewing a resiliency plan	108
	Creating a schedule for a resiliency plan	109
	Editing a schedule for a resiliency plan	109
	Deleting a schedule for a resiliency plan	110
	Viewing a schedule for a resiliency plan	110
Appendix A	General troubleshooting	112
	Viewing events and logs in the console	112
	Events in VMware virtual machines disaster discovery	113
	Troubleshooting discovery of assets	113
	Log files to troubleshoot Veritas Resiliency Platform Data Mover	116
	Resiliency Platform fails to attach storage policy to virtual machines	117
	Resiliency Platform fails to create storage policy	117
	Resolving the Admin Wait state	117
	Admin Wait state codes	118
Glossary		122

Index	124
-------------	-----

Overview of Resiliency Platform

This chapter includes the following topics:

- [About Veritas Resiliency Platform](#)
- [About disaster recovery using Resiliency Platform](#)
- [About Resiliency Platform features and components](#)
- [About Resiliency Platform capabilities](#)
- [About managing VMware virtual machines using Resiliency Platform](#)
- [About permissions for operations in the console](#)

About Veritas Resiliency Platform

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

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

Resiliency Platform has the following core capabilities:

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

About disaster recovery using Resiliency Platform

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

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

- Monitoring of data center assets - storage, virtual machines, and applications.
- Ability to group your virtual machines or applications in resiliency groups based on your production environment and business needs.
- Making business services more resilient by providing the ability to perform disaster recovery operations on virtual machines and applications. For example, migrate and take over.
- Ability to replicate data from virtual machines on source data centers to target data centers using Resiliency Platform Data Mover integrated with VMware API I/O filtering framework or array-based replication technologies provided by array vendors.
- Resiliency plan (a sequential execution of predefined steps) to automate site-level recovery operations on your IT infrastructure in the event of downtime.

- Auto-discovery and real-time tracking for recovery objectives.
- Ability to perform non-disruptive testing (rehearsal) on your virtual machines and applications to ensure that your infrastructure is adequately prepared for protection in the event of disaster.
- Reporting capabilities providing details about resiliency health of applications and virtual machines.

See [“Understanding the role of resiliency groups in disaster recovery operations”](#) on page 65.

About Resiliency Platform features and components

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

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

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

About Resiliency Platform capabilities

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

Table 1-1 Resiliency Platform capabilities

Capability	More information
Configuring virtual machines and applications for remote recovery operations or basic monitoring	See “Managing virtual machines for remote recovery (DR)” on page 67. See “Managing virtual machines for basic monitoring” on page 18.
Starting and stopping resiliency groups for maintenance	See “Starting a resiliency group” on page 19. See “Stopping a resiliency group” on page 20.
Rehearsing disaster recovery	See “Performing the rehearsal operation” on page 76.
Migrating a resiliency group	See “Migrating a resiliency group of virtual machines” on page 80.
Taking over resiliency groups	See “Taking over a resiliency group of virtual machines” on page 82.
Performing the resync operation	See “Performing the resync operation” on page 82.
Managing activities and resiliency plans	See “Managing activities” on page 98. See “Managing resiliency plans” on page 100.
Displaying an overview of your resiliency domain including the number and health of your resiliency groups	See “About the Resiliency Platform Dashboard” on page 84. See “Displaying resiliency group information and status” on page 21.
Monitoring risks for protected assets	See “About risk insight” on page 89.
Viewing reports	See “Viewing reports” on page 87.

About managing VMware virtual machines using Resiliency Platform

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

Note: Ensure that you install VMware tools for VMware virtual machines.

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

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

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

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

Disaster recovery configuration requires that you set up replication for your virtual machines. You can set up replication using storage arrays or you can use Resiliency Platform Data Mover.

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

About permissions for operations in the console

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

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

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

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

Managing resiliency groups

This chapter includes the following topics:

- [About resiliency groups](#)
- [Guidelines for organizing resiliency groups](#)
- [About service objectives](#)
- [Managing virtual machines for basic monitoring](#)
- [Starting a resiliency group](#)
- [Stopping a resiliency group](#)
- [Displaying resiliency group information and status](#)
- [Viewing resiliency group details](#)
- [Editing a resiliency group](#)
- [Deleting a resiliency group](#)

About resiliency groups

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

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

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

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

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

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

See [“About service objectives”](#) on page 17.

See [“Managing virtual machines for basic monitoring”](#) on page 18.

See [“Managing virtual machines for remote recovery \(DR\)”](#) on page 67.

Guidelines for organizing resiliency groups

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

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

- Ensure that all the virtual machines that are to be grouped in a single resiliency group are from a single hypervisor or virtualization server (if not clustered) or a single cluster.
- Ensure that they consume storage from the same Veritas Replication Set. E.g. EMC SRDF device group, NetApp Volume, 3PAR replication group, and so on.

About service objectives

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

Such service objectives can include tunables such as Recovery Point Objective (RPO) for assets being managed in that resiliency group.

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

For example, if you apply a recovery service objective that supports recovery of virtual machines to a remote data center, you must select the recovery data center.

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

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

Note: Authorization to perform operations also depends on the permissions assigned to users and groups in Resiliency Platform.

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

Managing virtual machines for basic monitoring

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

There are two types of pre-activated service objectives:

- Monitor assets - provides only monitoring, start, and stop operations
- Recover hosts - provides recovery operations as well as the start and stop operations

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

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

See [“Managing virtual machines for remote recovery \(DR\)”](#) on page 67.

To manage virtual machines for basic monitoring

1 Prerequisites

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

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

2 Navigate



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

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

3 Select the virtual machines:

- Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.
- Drag and drop virtual machines to **Selected Instances**.
For VMware assets, if you plan to configure disaster recovery protection later using Resiliency Platform Data Mover, each resiliency group must map to only one ESX cluster.

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

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

6 Supply a name for the resiliency group.

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

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

Starting a resiliency group

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

To start a resiliency group

1 Navigate



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

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

3



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

You can also perform operations from the Details page.

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

If you have applied update 2.0.0.100 on Veritas Resiliency Platform 2.0, you can select the checkbox on the **Start Resiliency Group** wizard to start the post-replication operations of migrate or takeover workflow on the production data center such as refreshing storage, network, compute, and customization.

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

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

Stopping a resiliency group

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

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

To stop a resiliency group


1 Prerequisites

- Make sure that you are aware of all the assets in the resiliency group, and the potential effect on users if you shut them down.
- Choose a time for stopping the resiliency group that minimizes any disruption of service.
- If necessary, notify users before you stop the resiliency group.

2 Navigate



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

- 3 Locate the resiliency group. Use filters or Search as needed.
- 4  On the row for the resiliency group, select the vertical ellipsis > **Stop**.
You can also perform operations from the Details page.
- 5 On the **Stop Resiliency Group** screen, select the data center in which to stop the resiliency group, and submit.

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

Displaying resiliency group information and status

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

Table 2-1 Displaying resiliency group information and status


Location	Level of detail	Useful for
Resiliency Platform Dashboard	Lowest. Displays the number of resiliency groups under Resiliency Platform control and the total number of groups in error, at risk, and healthy.	Getting a quick overview of the resiliency group population and health throughout Resiliency Platform. See “About the Resiliency Platform Dashboard” on page 84.
 Assets > Resiliency Groups tab	Medium. Lists all your resiliency groups in one place.	Seeing what is in each of your data centers, the state of the groups, and so on.

Table 2-1 Displaying resiliency group information and status *(continued)*

Location	Level of detail	Useful for
Resiliency group-specific screen	Highest. Lists each asset in the resiliency group, their type, and state.	Getting detailed information on a resiliency group and its underlying assets, including disaster recovery status. This screen lists available operations for the group. See “Viewing resiliency group details” on page 24.

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

To display resiliency group information and status

1 Navigate



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

2 Review information and status

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

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

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

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

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

Viewing resiliency group details

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

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

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

To view details of a resiliency group

1 Navigate



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

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

3 On the row for the resiliency group, select the vertical ellipsis > **Details**. You can also double-click the row to view details.

The details page includes the following:

- Menu options for operations that you can perform on the resiliency group.
- Details of how the resiliency group is configured.
- Status information.
- A list of the resiliency group assets and their state.

See [“Displaying resiliency group information and status”](#) on page 21.

Editing a resiliency group

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

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

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

To edit the resiliency group information

1 Prerequisites

- Determine the potential impact editing the resiliency group may have on users. If necessary, notify users of the upcoming change.

2 Navigate



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

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

- #### 4
- On the row for the resiliency group, select the vertical ellipsis > **Edit**.
You can also edit the resiliency group from its Details page.

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

Deleting a resiliency group

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

To successfully complete the delete operation ensure the following:

- The assets on the production data center are running and accessible.
- The xprtld daemon on the virtual machines is running.

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

- During the operation, replication is stopped and Veritas Replication Sets are deleted on gateways and on-premises virtual machines.
- Journal disks are removed from the virtual machines on the production data center and cloud virtual machines instances are deleted. For physical machines, the journal disks are unsigned.

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

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

To delete a resiliency group

1 Navigate



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

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

3 On the row for the resiliency group, select the vertical ellipsis > **Delete**. You can also perform operations from the Details page

4 Confirm the deletion.

Overview of Resiliency Platform Data Mover

This chapter includes the following topics:

- [About Veritas Resiliency Platform Data Mover](#)
- [Supported environments for Resiliency Platform Data Mover with VMware VAIO](#)
- [Resiliency Platform Data Mover deployment process overview](#)
- [How Resiliency Platform Data Mover works](#)
- [About Veritas Resiliency Platform Data Mover architecture](#)
- [About full synchronization with Veritas Resiliency Platform Data Mover](#)
- [How Veritas Resiliency Platform Data Mover handles virtual machine writes](#)
- [How virtual appliances for Veritas Resiliency Platform Data Mover are deployed and configured](#)

About Veritas Resiliency Platform Data Mover

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

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

Veritas Resiliency Platform Data Mover allows replication of only VMware virtual machines. Veritas Resiliency Platform Data Mover provides data replication between

geographically separated data centers facilitating an effective disaster recovery solution.

Features of Veritas Resiliency Platform Data Mover include the following:

- Replicates virtual machines including its boot and data disks from source data center to target data center over any IP network in a LAN or a WAN environment.
- Enables easy recovery of virtual machines in the target data center.
- Ensures virtual machine data consistency.
- Recovers virtual machines protected by Data Mover at the Resiliency Group level.
- Enables non-disruptive testing of recovery at target data centers.

Supported environments for Resiliency Platform Data Mover with VMware VAIO

Veritas Resiliency Platform Data Mover is supported on VMware clusters, with ESXi 6.0 U2 and above versions.

- The Veritas Resiliency Platform Data Mover virtual appliance must be part of the same ESX cluster in source and target data centers.
- Each host system in source and target data center must run ESXi version 6.0 U2 or later versions.
- Check the acceptance level on every host system in source and target data center. Ensure that the host system's acceptance level is not set to `VMwareCertified`.

The allowed acceptance levels are `VMwareAccepted`, `PartnerSupported`, or `CommunitySupported`. By default, the ESX host is set to `PartnerSupported`.

- VMware tools are installed on virtual machines.
- Virtual machine disks must be in the same datastore.
- Virtual machines must not have any physical RDM disks.

Resiliency Platform Data Mover deployment process overview

The following is a summary of the steps that are required to implement Resiliency Platform Data Mover and where to go for more information on each step.

Table 3-1 Resiliency Platform Data Mover deployment process overview

Step	More information
Deploy the Replication Gateway virtual appliances on the production and the recovery data centers	<i>Veritas Resiliency Platform Deployment Guide</i> See “How virtual appliances for Veritas Resiliency Platform Data Mover are deployed and configured” on page 34.
Configure the Replication Gateway virtual appliances	<i>Veritas Resiliency Platform Deployment Guide</i>
Configure the Replication Gateway pair	See “Creating a Veritas Replication Gateway pair” on page 61.
Add the vCenter Server and the discovered ESX server assets to Resiliency Platform at the data centers	<i>Veritas Resiliency Platform Deployment Guide</i>
Configure resiliency groups of virtual machines for disaster recovery with Resiliency Platform Data Mover technology	See “Managing virtual machines for remote recovery (DR)” on page 67.

How Resiliency Platform Data Mover works

The Veritas Resiliency Platform Data Mover feature of Resiliency Platform replicates all the virtual machine writes at the local (source) data center to the remote (target) data center. The replication provides a consistent copy of the data. If a disaster occurs at the source data center, Resiliency Platform can use the copy of the data on the target (remote) data center to provision and start a virtual machine on the remote data center.

To protect virtual machines using Resiliency Platform Data Mover, you group the virtual machines at the source data center into resiliency groups that use Resiliency Platform Data Mover to provide disaster recovery protection. The resiliency group is the unit of recovery, so the virtual machines that need to be recovered together must be in the same resiliency group.

During the configuration process, Resiliency Platform puts virtual machines into multiple Veritas Replication Sets and replication units associated to the virtual machine. Each Veritas Replication Set caters to a single virtual machine and includes all the disks attached to that virtual machine, including boot and data disks. Each constituent disk is referred to as a Replication Unit.

When an application or virtual machine runs, several processes perform writes to disks, in a specific order. For example, a database posts any database change to the log before writing to the table space. The term write-order fidelity means that

the write order across the constituent disks or replication units is maintained at all times.

Resiliency Platform Data Mover maintains write-order fidelity for a Veritas Replication Set when the replication is in the active state. The write-order fidelity ensures that the data in the target data center is consistent. Even though data at the target data center may not be the most recent copy, Data Mover makes sure that this data is always consistent.

Resiliency Platform Data Mover tracks writes for the virtual machines on the source data center in the order in which they are received. It applies the writes on the target data center in the same order, thereby maintaining write order fidelity.

The replication includes any changes to the boot disks of the virtual machines. As a result, if a disaster occurs on the source data center, or a planned migration is performed, virtual machines can be brought up on the recovery data center. The disaster recovery operation in Resiliency Platform provisions the virtual machines in the recovery data center so that they can be brought online as part of the operation.

About Veritas Resiliency Platform Data Mover architecture

Veritas Resiliency Platform Data Mover uses the VMware vSphere APIs for I/O Filtering (VAIO) framework. It provides an I/O filter for replication that has been certified by VMware. VAIO enables filtering of a virtual machine's I/Os. The filter runs inside an ESXi server and intercepts I/O requests moving between a guest operating system and virtual disks.

Veritas Resiliency Platform Data Mover deploys a user-land module, called **vtstap**, on the ESXi host where the protected VMs are running. This module is built using the VMware VAIO APIs. It intercepts and replicates I/Os from the virtual machines. When you configure disaster recovery with Resiliency Platform Data Mover, the configuration process installs this module on the ESXi hosts for the selected virtual machines on the source data center. It also installs this module on the ESXi hosts in the selected cluster on the target data center. After a migrate or takeover operation is performed, the platform ensures that data is available in the target data center.

After Veritas Resiliency Platform Data Mover filters I/Os in the user land of the ESXi host where virtual machines are running, the I/O goes through the following path within the Replication Gateway at the source data center before it is replicated to the target data center.

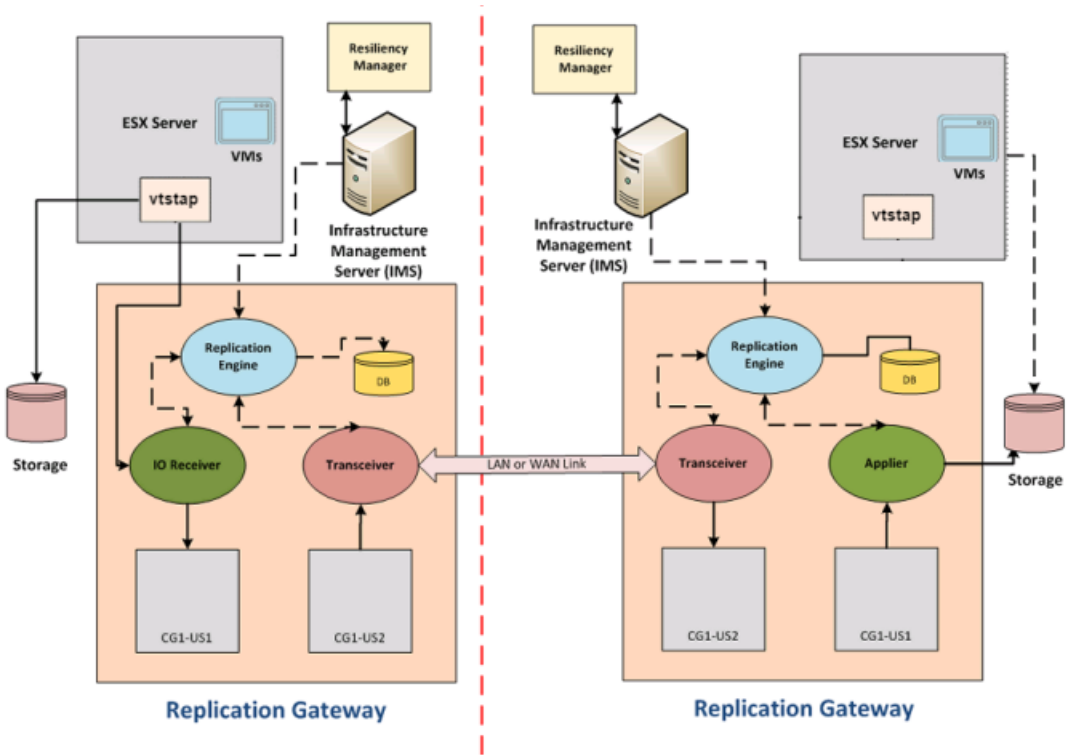
Each Replication Gateway includes four daemons that run when replication is enabled:

- **I/O receiver**
 Continuously receives the virtual machine I/Os that were tapped and sent by the vtstap module in the ESXi host where virtual machines are running in a continuous fashion.
- **Transceiver**
 Transfers and receives data over the WAN link periodically.
- **Applier**
 Applies the data to the storage after it is received on the target gateway.
- **Scheduler**
 Manages and schedules data transfer between gateways.
- **Engine**
 Maintains the state of replication and also coordinates with all other components.

The virtual machines on the target (recovery) data center are provisioned only when a disaster recovery operation (such as migrate) is run in Resiliency Platform. The disaster recovery operation then can bring the virtual machines online in the recovery data center. This avoids unnecessary resource utilization and accounting when the workload is running in the other data center.

To use Veritas Resiliency Platform Data Mover, the source Replication Gateway and the target Replication Gateway are linked together into a Replication Gateway pair, which establishes the replication channel between the source and the target. A Replication Gateway pair is a one-to-one mapping of the source Replication Gateway to the target Replication Gateway. You can choose to encrypt the communication between gateways, unless you are using a dedicated VPN link.

Figure 3-1 Replication architecture



About full synchronization with Veritas Resiliency Platform Data Mover

When Data Mover is configured for a resiliency group, replication is started. At that time, the storage on the target data center must be synchronized with the data from the source data center. This process of synchronizing the entire set of data is known as full synchronization.

The amount of time that is required for full synchronization depends on several factors. These factors include the size of the replication disks, the network and resiliency of the LAN and WAN environment, and the amount of IO occurring during the synchronization. After the full synchronization is complete, the replication moves into active state. In the active state, Data Mover maintains write-order fidelity.

A full synchronization is also required after a takeover operation to prepare for the next takeover or migration. Takeover is an activity initiated by a user when the source data center is down due to a disaster, and the virtual machines need to be

brought up at the target (recovery) data center to provide business continuity. After a takeover, the virtual machine runs in the target (recovery) data center. Once the source (production) data center is back up and running, you must perform a Resync operation from the recovery data center before you can migrate back to the production data center. The Resync operation synchronizes the data on the production data center with the data in the recovery data center storage. When the synchronization completes, the production data center is up-to-date. You can then perform the Migrate operation.

How Veritas Resiliency Platform Data Mover handles virtual machine writes

Veritas Resiliency Platform Data Mover uses the vtstap service running on ESX host, to intercept and process protected virtual machine writes. The vtstap service intercepts the writes to the storage, while reads are directly processed from the virtual machine storage.

The vtstap service records the location of the write I/O, sequences the writes, and writes are applied to the virtual machine storage. Sequenced writes are then asynchronously sent to source replication gateway. The writes that accumulate on the source Replication Gateway are periodically sent to the target Replication Gateway. The target Replication Gateway applies the writes to the target data center storage. This ensures that data is consistent on the source and target data centers. As Resiliency Platform Data Mover employs asynchronous replication, there might be lag between source and target, but it will always be consistent.

Resiliency Platform Data Mover processes an incoming write by performing the following steps in the order listed:

- The operating system in the guest VM issues a write to the virtual machine storage.
- The I/O tap module (vtstap) records the location of the I/O.
- IO is written to virtual machine storage.
- The vtstap module sends the I/O data over the network to the I/O receiver in the source Replication Gateway.
- The I/O receiver aggregates the I/Os.
- Periodically, the aggregated I/Os are sent to the transceiver.
- The transceiver sends the I/Os across the network to the transceiver on the target Replication Gateway.

- The I/O is sent to the applier once the transceiver on the target replication gateway receives the set of I/Os.
- The applier writes the I/O to the target data center storage.
- The operating system in the guest VM issues a write to the virtual machine storage.
- The I/O tap module (vtstap) records the location of the I/O.
- IO is written to virtual machine storage.
- The vtstap module sends the I/O data over the network to the I/O receiver in the source Replication Gateway.
- The I/O receiver aggregates the I/Os.
- Periodically, the aggregated I/Os are sent to the transceiver.
- The transceiver sends the I/Os across the network to the transceiver on the target Replication Gateway.
- The I/O is sent to the applier once the transceiver on the target replication gateway receives the set of I/Os.
- The applier writes the I/O to the target data center storage.

How virtual appliances for Veritas Resiliency Platform Data Mover are deployed and configured

Veritas Resiliency Platform includes a virtual appliance for Veritas Resiliency Platform Data Mover, which is configured as the Veritas Replication Gateway. You deploy at least two Replication Gateway virtual appliances. One Replication Gateway is deployed in the source data center and the other Replication Gateway is deployed in the target data center.

The virtual appliance for the Replication Gateway is deployed using the same process as the virtual appliance for the other Resiliency Platform components.

However, the Replication Gateway on the source data center must have access to the datastores where protected virtual machines reside. The Replication Gateway on the target data center must have access to the datastores where recovery virtual machines would be provisioned.

After deploying the Replication Gateway as a virtual appliance, bootstrap appliance to configure the Replication Gateway with network and other settings. During the product bootstrap configuration, you also register the gateway with an Infrastructure Management Server (IMS) on the same data center .

How virtual appliances for Veritas Resiliency Platform Data Mover are deployed and configured

From that point, the additional tasks that are required to use Resiliency Platform Data Mover to protect virtual machines are done from the Resiliency Platform web console.

Preparing for disaster recovery configuration

This chapter includes the following topics:

- [Supported replication technologies with Veritas Resiliency Platform](#)
- [About the Veritas Resiliency Platform Data Mover Replication Gateway](#)
- [DR for VMware virtual machines using Resiliency Platform Data Mover - an overview of key steps](#)
- [DR for VMware virtual machines using array-based replication - an overview of key steps](#)
- [Prerequisites for configuring VMware virtual machines for disaster recovery](#)
- [VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover](#)
- [Limitations for virtual machine disaster recovery](#)

Supported replication technologies with Veritas Resiliency Platform

When you configure virtual machines for disaster recovery, Veritas Resiliency Platform lets you select the replication technology to replicate data from a source data center to a target data center.

Veritas Resiliency Platform supports the following replication technologies. Depending on your environment, select the replication technology that best fits your business needs.

- Veritas Resiliency Platform Data Mover replication, which is integrated with VMware vSphere API I/O filtering (VAIO) framework
- Array-based replication technologies that are provided by the following array vendors: EMC SRDF, EMC Recoverpoint, Netapp (cDOT) Snapmirror, HP 3PAR Remote Copy, Hitachi TrueCopy/HUR

About the Veritas Resiliency Platform Data Mover Replication Gateway

The Resiliency Platform Data Mover feature of Veritas Resiliency Platform includes a virtual appliance, called Veritas Replication Gateway that can be deployed for replication. The gateway optimizes data replication by aggregating and batching data from multiple virtual machines. It further, write cancels data and batches data before transferring it to the target data center.

Each Replication Gateway must be paired with a Replication Gateway at another data center before replication is enabled.

During the process of deploying and configuring a gateway virtual appliance, you register it with an Infrastructure Management Server (IMS) on that data center.

DR for VMware virtual machines using Resiliency Platform Data Mover - an overview of key steps

To configure virtual machines with Resiliency Platform Data Mover ensure that virtual machines are not already configured with another replication technology. Before you can perform disaster recovery (DR) operations such as rehearse, migrate, and takeover on virtual machines using Resiliency Platform Data Mover, you must prepare the replication environment, and configure assets for disaster recovery. This section lists the key steps required to configure the disaster recovery of VMware virtual machines using Veritas Resiliency Platform Data Mover.

Table 4-1 Configuring disaster recovery using Data Mover

Action	Description	Refer to
Deploy the Replication Gateways	Deploy and configure the Veritas Replication Gateway virtual appliances on both data centers	Refer to the Deployment Guide.
Create the gateway pairs	Configure the Replication Gateway pairs to be used for replication	See “Creating a Veritas Replication Gateway pair” on page 61.
Add the VMware asset infrastructure	Prepare the VMware virtualization servers for the Resiliency Platform environment Add the VMware virtualization servers to Resiliency Platform at both data centers	Refer to the Deployment Guide.
Prepare the virtual machines	Ensure that VMware Tools and other prerequisites are configured on virtual machines	Refer to the Deployment Guide. See “Prerequisites for configuring VMware virtual machines for disaster recovery” on page 50. See “Limitations for virtual machine disaster recovery” on page 54.

Table 4-1 Configuring disaster recovery using Data Mover *(continued)*

Action	Description	Refer to
Storage requirements for Replication Gateway	Both the source and recovery (target) gateway must have external storage equivalent to 12GB for each virtual machine protected by the gateway pair. A maximum of 58 volumes or disks can be attached to the Replication Gateway.	See “Prerequisites for configuring VMware virtual machines for disaster recovery” on page 50.
Configure your assets for disaster recovery	Group the required virtual machines in a resiliency group and choose the appropriate service objective to configure disaster recovery for the resiliency group	See “Managing virtual machines for remote recovery (DR)” on page 67.
Rehearse DR operations	Test your disaster recovery environment to ensure readiness	See “Performing the rehearsal operation” on page 76. See “Performing cleanup rehearsal ” on page 77.

DR for VMware virtual machines using array-based replication - an overview of key steps

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

Table 4-2 Configuring disaster recovery using array-based replication

Action	Description	Refer to
Set up your replication environment	Set up your VMware environment and storage arrays for replication	<p>See “Configuring VMware virtual machines disaster recovery using EMC SRDF replication” on page 41.</p> <p>See “Configuring VMware disaster recovery using NetApp SnapMirror” on page 43.</p> <p>See “Configuring VMware virtual machines disaster recovery using EMC RecoverPoint replication” on page 45.</p> <p>See “Configuring VMware virtual machines disaster recovery using Hitachi True Copy replication” on page 47.</p> <p>See “Configuring VMware virtual machines disaster recovery using HPE 3PAR Remote Copy replication” on page 48.</p>
Add the asset infrastructure	Add the VMware servers and the storage arrays used for replication to Resiliency Platform	Refer to the Deployment Guide.
Prepare the virtual machines	Ensure that VMware Tools and other prerequisites are configured on virtual machines	<p>Refer to the Deployment Guide.</p> <p>See “Prerequisites for configuring VMware virtual machines for disaster recovery” on page 50.</p> <p>See “Limitations for virtual machine disaster recovery” on page 54.</p>
Configure network settings	Configure network settings for mapping between data centers	See “Setting up network mapping between production and recovery data centers” on page 58.
Configure your assets for disaster recovery	Group the virtual machines in a resiliency group and apply the appropriate service definition	See “Managing virtual machines for remote recovery (DR)” on page 67.
Rehearse DR operations	Test your disaster recovery environment to ensure readiness	See “Performing the rehearsal operation” on page 76.

Configuring VMware virtual machines disaster recovery using EMC SRDF replication

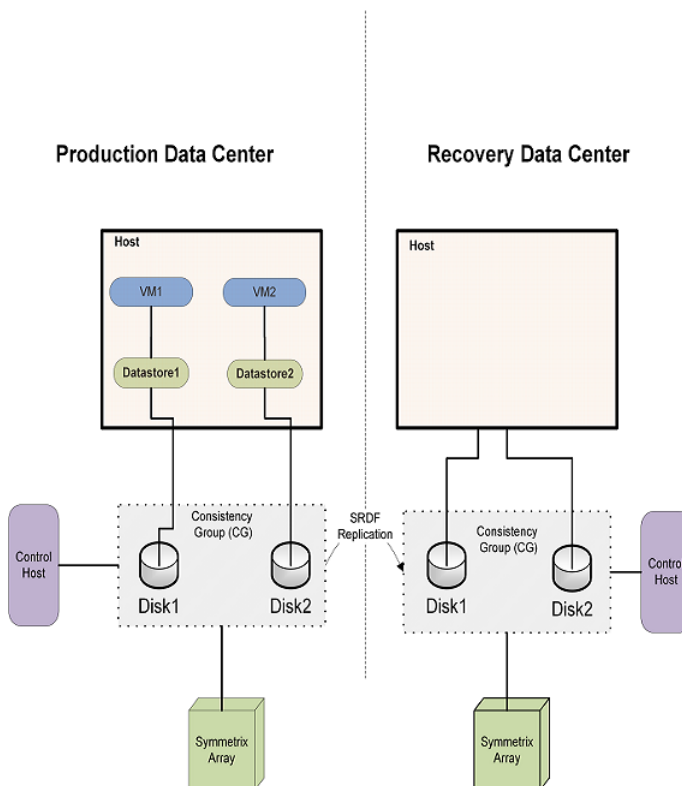
This section lists the prerequisites to enable data replication using EMC SRDF for the Veritas Resiliency Platform environment.

- Ensure that EMC Solutions Enabler is installed on a host and that the SRDF device groups are already set up for the replication between the primary and remote arrays.
- Ensure that the SRDF replicated LUNs are assigned to the respective VMware ESX Servers. Do not attach replicated peer SRDF LUNs (R1 and R2) to the same VMware ESX Server.

Note: If any changes are made to storage or storage-adapters available to the ESX/ESXi host or cluster, you must perform storage rescan, storage adapter rescan or both depending on the changes made.

- Ensure that the virtual machine's virtual disk files that are replicated using EMC SRDF replication are stored on a datastore and a corresponding device group must be created for them.

Note: For EMC SRDF-based replication in Resiliency Platform, all virtual machines that consume storage from a Veritas Replication Set must belong to the same resiliency group.



Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks to add the asset infrastructure to Resiliency Platform for discovery by the Infrastructure Management Server (IMS).

Resiliency Platform configurations:

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

See the *Veritas Resiliency Platform Deployment Guide*.

- Add the Symmetrix enclosures to the appropriate data centers. Provide the discovery host name and the SYMCLI location on this discovery host.

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

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

Note: Any managed host can be designated as the array discovery host, including the virtual machine inside VMware ESX server that has EMC Symmetrix Gatekeeper device visibility and SYMCLI installed. The host on which consistency groups are defined can also be used as an array discovery host.

This operation returns the list of Symmetrix arrays (local and remote) accessible to the discovery host. To configure disaster recovery for the virtual machines, select one or more local arrays only.

Ensure that the enclosure discovery is complete before proceeding with adding the VMware vCenter Servers

- Add the vCenter Servers to the appropriate data centers. The user needs to have vCenter administrator privileges.

Note: Ensure that the virtualization server and ESX server are discovered successfully.

- Add the host where the SRDF device groups are configured to the appropriate data centers.

See [“Troubleshooting discovery of assets”](#) on page 113.

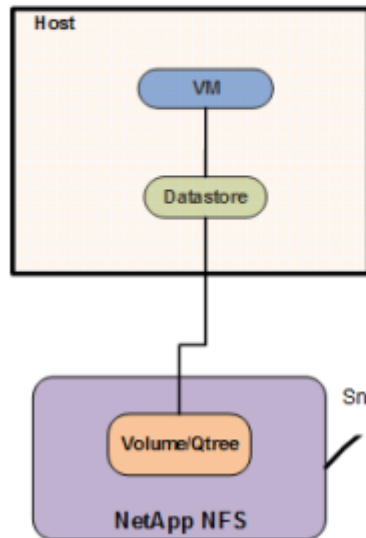
Configuring VMware disaster recovery using NetApp SnapMirror

This section lists the prerequisites to enable data replication for the Veritas Resiliency Platform environment using NetApp SnapMirror. For NetApp SnapMirror based replication all virtual machines that consume storage from a NetApp volume must belong to the same resiliency group.


- Ensure that the NetApp volumes are already set up for replication between the primary and remote NetApp storage systems, and the replication has a replication schedule associated with it.
Resiliency Platform does not support one NetApp volume having more than one SnapMirror destination volumes.
- NetApp volumes or qtrees can be provisioned as NFS datastores mounted on ESX servers, or LUNs created on NetApp volumes or qtrees can be provisioned to ESX servers via Fibre Channel and mounted as VMFS datastores.

- Ensure that the NetApp SnapMirror replicated volumes are mounted on the respective VMware ESX servers in both the sites. Do not mount the replicated peer NetApp volumes to the same VMware ESX server.
Also ensure that the volumes are replicated using SnapMirror policy type mirror. SnapMirror policy types vault and mirror-vault are not supported.
- If any changes are made to storage or storage-adapters available to the ESX/ESXi host or cluster, you must perform storage rescan, storage adapter rescan or both depending on the changes made.

Production Data Center



Recovery Data Center



SnapMirror based
replication

Resiliency Platform configurations:

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

See the *Veritas Resiliency Platform Deployment Guide*.

Resiliency Platform configurations:

- Configure the VMware vCenter Server to send traps to the IMS.
- Add the VMware vCenter Servers to their respective data centers. The user needs to have vCenter administrator privileges.

Note: Ensure that the virtualization server and ESX server are discovered successfully.

- Add the NetApp enclosures to their respective data centers.
While adding the enclosure, you must provide the IMS name as the discovery host name. Also provide the NetApp storage system name or IP, and credentials. The user should have sufficient privileges to perform SnapMirror replication operations.

See [“Troubleshooting discovery of assets”](#) on page 113.

Configuring VMware virtual machines disaster recovery using EMC RecoverPoint replication

This section lists the prerequisites to enable data replication using EMC RecoverPoint for the Veritas Resiliency Platform environment.

- Ensure that discovery host is able to communicate with RecoverPoint appliance using SSH.
- Ensure that the RecoverPoint replicated LUNs are assigned to the respective VMware ESX Servers. Do not attach replicated peer RecoverPoint LUNs (Production and remote copy) to the same VMware ESX Server.

Note: If any changes are made to storage or storage-adapters available to the ESX/ESXi host or cluster, you must perform storage rescan, storage adapter rescan or both depending on the changes made.

- Ensure that the virtual machine's virtual disk files that are replicated using EMC RecoverPoint replication are stored on a datastore and a corresponding device group must be created for them.

Note: For EMC RecoverPoint-based replication in Resiliency Platform, all virtual machines that consume storage from a Veritas Replication Set must belong to the same resiliency group.

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks to add the asset infrastructure to the Infrastructure Management Server (IMS).

Resiliency Platform configurations:

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

See the *Veritas Resiliency Platform Deployment Guide*.

- Add the backend enclosures to the appropriate data centers. Provide the discovery host name and the SYMCLI location on this discovery host.

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

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

Note: Any managed host can be designated as the array discovery host, including the virtual machine inside VMware ESX server that has EMC Symmetrix Gatekeeper device visibility and SYMCLI installed. The host on which consistency groups are defined can also be used as an array discovery host.

This operation returns the list of Symmetrix arrays (local and remote) accessible to the discovery host. To configure disaster recovery for the virtual machines, select one or more local arrays only.

Ensure that the enclosure discovery is complete before proceeding with adding the VMware vCenter Servers

- Add the vCenter Servers to the appropriate data centers. The user needs to have vCenter administrator privileges.

Note: Ensure that the virtualization server and ESX server are discovered successfully.

See [“Troubleshooting discovery of assets”](#) on page 113.

Configuring VMware virtual machines disaster recovery using Hitachi True Copy replication

This section lists the prerequisites to enable data replication using Hitachi True Copy (HTC) for the Veritas Resiliency Platform environment.

- Ensure that Hitachi Command Control Interface (HORCM CCI) is installed on a host and that the HTC Instances are already set up for the replication between the primary and remote arrays.
- Ensure that HORCM CLI executes properly on the host. This is required for discovery and all other operations.
- Ensure that the HTC replicated LUNs are assigned to the respective VMware ESX Servers.

Note: If any changes are made to storage or storage-adapters available to the ESX/ESXi host or cluster, you must perform storage rescan, storage adapter rescan or both depending on the changes made.

- Ensure that the virtual machine's virtual disk files that are replicated using Hitachi True Copy replication are stored on a datastore and a corresponding HTC Instances must be created for them.

Note: For HTC based replication in Resiliency Platform, all virtual machines that consume storage from an HTC Instance must belong to the same resiliency group. An Instance is a collection of Volume Groups that helps in maintaining write consistency during replication. A resiliency group is a unit of management and control that you create in Resiliency Platform. Related virtual machines are organized into a resiliency group and managed and monitored as a single entity.

Resiliency Platform configurations:

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

See the *Veritas Resiliency Platform Deployment Guide*.

Resiliency Platform configurations:

- Configure the VMware vCenter Server to send traps to the IMS.
- Add the HTC discovery hosts to the respective data centers. HTC discovery hosts are hosts on which HORCM package is installed and HTC device groups are created.

- Add the vCenter Servers to their respective data centers. The user needs to have vCenter administrator privileges.

Note: Ensure that the virtualization server and ESX server are discovered successfully.

- Add the host where the HTC Instances are configured to their respective data centers.

Configuring VMware virtual machines disaster recovery using HPE 3PAR Remote Copy replication

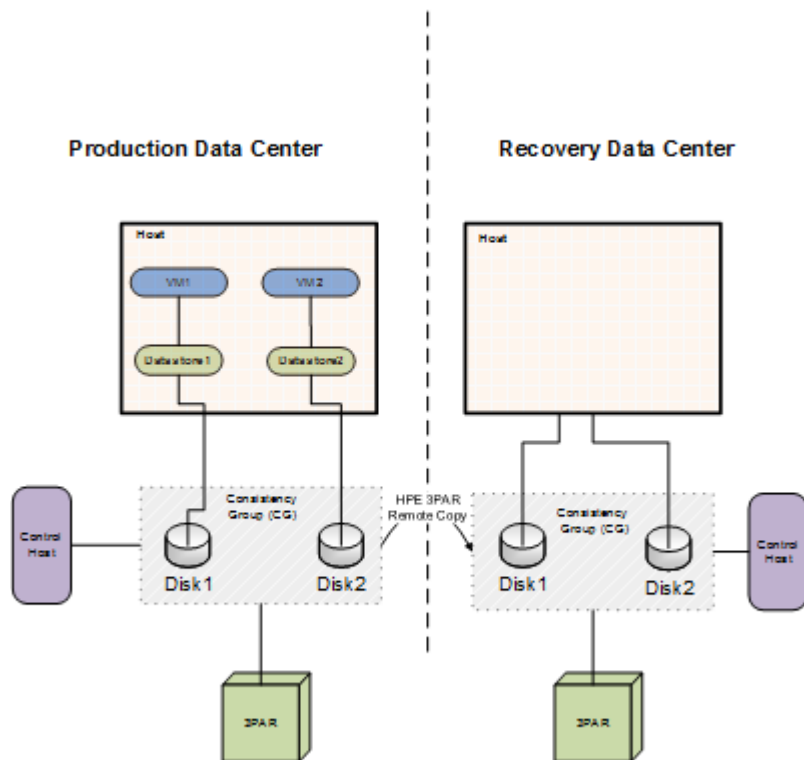
This section lists the prerequisites to enable data replication using HPE 3PAR Remote Copy replication for the Veritas Resiliency Platform environment.

- Ensure that the discovery host is able to communicate with HPE 3PAR array using SSH.
- Confirm that HPE 3PAR array user has all the necessary permissions to perform HPE 3PAR RemoteCopy operations.
- HPE Remote Copy groups are set up for replication between the primary and the remote arrays. Ensure that the group names are unique across all data centers. Group name on the recover data center is auto generated by HPE. Do not modify the name.
- Ensure that the HPE 3PAR Remote Copy replicated LUNs are assigned to the respective VMware ESX Servers.

Note: If any changes are made to storage or storage-adapters available to the ESX/ESXi host or cluster, you must perform storage rescan, storage adapter rescan or both depending on the changes made.

- Ensure that the virtual machine's virtual disk files that are replicated using HPE 3PAR Remote Copy replication are stored on a datastore and a corresponding 3PAR Remote Copy instances must be created for them.

Note: For HPE 3PAR Remote Copy based replication in Resiliency Platform, all virtual machines that consume storage from a 3PAR Remote Copy instance must belong to the same resiliency group. An Instance is a collection of Volume Groups that helps in maintaining write consistency during replication. A resiliency group is a unit of management and control that you create in Resiliency Platform. Related virtual machines are organized into a resiliency group and managed and monitored as a single entity.



Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks to add the asset infrastructure to the Infrastructure Management Server (IMS). The following is a summary of the steps. More information is available.

See the *Veritas Resiliency Platform Deployment Guide*.

Resiliency Platform configurations:

- Configure the VMware vCenter Server to send traps to the IMS.
- Add the 3PAR enclosure to the IMS using the **+ HP Enclosure** option.

Note: Any managed host can be designated as the array discovery host, including the virtual machine inside VMware ESX server.

This operation returns the list of 3PAR arrays (local and remote) accessible to the discovery host. To configure disaster recovery for the virtual machines, select one or more local arrays only. Ensure that the enclosure discovery is complete before proceeding with adding the VMware vCenter Servers.

- Add the vCenter Servers to their respective IMS in each data center using the **+ vCenter** option. The user needs to have vCenter administrator privileges.

Note: Ensure that the virtualization server and ESX server are discovered successfully.

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

Limitations:

- Raw Device Mapping (RDM) on replicated disks to virtual machine is not supported.
- HPE 3PAR Remote Copy synchronous replication is not supported.
- 3PAR storage connectivity via iSCSI is not supported.

Prerequisites for configuring VMware virtual machines for disaster recovery

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:

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

- 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. Also all the virtual disks must belong to a single datastore.

Additional prerequisites for the virtualization environment depend on the type of replication you are using.

- For array-based replication, the replicated storage must be provisioned to the ESX servers at each data center.
- For Resiliency Platform Data Mover:
 - The Veritas Resiliency Platform Data Mover virtual appliance must be part of the ESX cluster in source and target data center.
 - Each host system in source and target data center must run ESXi version 6.0 U2 or later versions.
 - Ensure that no host system's acceptance level is `VMwareCertified`. The allowed acceptance level is `PartnerSupported`, which is the default level. Check the acceptance level on every host system in source and target data center.
 - VMware tools are installed on virtual machines.
 - Virtual machine disks should be in the same datastore.
 - Enough storage must be available on the recovery data center for provisioning the replicated virtual machines.
 - Both the source and recovery (target) gateway must have external storage equivalent to 12GB for each virtual machine protected by the gateway pair. For example, if a gateway pair supports 10 virtual machines, the source and recovery (target) gateway must each have 120GB of external storage.
 - A maximum of 58 volumes or disks can be attached to the Replication Gateway. The total number of disks from the virtual machines in the resiliency group and the number of disks that are already attached to the gateway should not exceed more than 58. This is a VMware limitation.
 - A virtual machine having snapshots cannot be protected using Resiliency Platform data Mover due to a VMware limitation.
 - The VMware vCenter Server user account used to add the server to Resiliency Platform must have the appropriate privileges for Data Mover. See [“VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover”](#) on page 52.

See [“DR for VMware virtual machines using Resiliency Platform Data Mover - an overview of key steps”](#) on page 37.

See [“DR for VMware virtual machines using array-based replication - an overview of key steps”](#) on page 39.

VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover

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

Table 4-3 VMware vCenter Server privileges required for Resiliency Platform Data Mover

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

Table 4-3 VMware vCenter Server privileges required for Resiliency Platform Data Mover *(continued)*

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

Limitations for virtual machine disaster recovery

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

Table 4-4

Limitations	Descriptions
Replication limitations	For more information on replication-based limitations of virtual machines, refer to the Hardware and Software Compatibility List (HSCL).
Limitations due to open-vm-tools on VMware virtual machines	<p>The guest IP reconfiguration operation which is required while performing the DR operations such as Migrate, Takeover, Rehearsals, is not supported by open-vm-tools.</p> <p>To reconfigure the guest IP you need to uninstall open-vm-tools and install the latest version of VMware Tools.</p> <p>Note that the virtual machine should not have both VMware Tools and open-vm-tools installed on it.</p> <p>Virtual machines with physical RDM disks cannot be replicated using Resiliency Platform Data Mover.</p>

Managing disaster recovery network mapping

This chapter includes the following topics:

- [Viewing and configuring network settings for a data center](#)
- [Editing network settings for a data center](#)
- [Removing network settings for a data center](#)
- [Configuring DNS server settings for a data center](#)
- [Setting up network mapping between production and recovery data centers](#)

Viewing and configuring network settings for a data center

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

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

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

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

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

To configure network settings for a data center

1 Navigate



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

Expand the data center > **Access Profile**

Click **+ Add Subnet**.

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

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

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

Editing network settings for a data center

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

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

To edit the network settings for a data center

1 Navigate



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

Expand the data center > **Access Profile**

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

2 Edit the name and the purpose as required.

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

Removing network settings for a data center

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

To remove the network settings for a data center

1 Navigate



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

Expand the data center > **Access Profile**

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

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

Configuring DNS server settings for a data center

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

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

To configure DNS server settings for a data center

1 Prerequisites

You must have the following information:

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

2 Navigate



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

Expand the data center > **Access Profile**

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

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

- 3 To add a new DNS server for the data center click **+ Add New DNS**.
- 4 Specify the IP address for the DNS server and select the purpose, either Rehearsal or Production.
- 5 Add one or more domains for the DNS server:
 - Fill in the domain name and the authentication type. For TSIG, browse to the key and private files. For GSSAPI, enter the user name and browse to the keytab file.
 - Enter a test host name and IP address and select **Test**. If the test is successful, that is the DNS configuration is validated, the **Add** button is enabled.
 - Select **Add**.
- 6 If you are done adding domains, select **Next**.
- 7 To remove a DNS server, right-click the required DNS server in the table and select **Remove**.

Setting up network mapping between production and recovery data centers

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

For VMware virtual machines, ensure that the mapping of all the concerned port groups (VLANs) across the data centers is configured before performing migrate, takeover, or rehearsal operations.

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

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

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

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

To set up network mapping between production and recovery data centers

1 Navigate



Disaster Recovery Settings (navigation pane)

Do one of the following:

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

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

- 2** In the **Network Mapping** page, select the source and the target data centers, and the network types that should be the part of your network pair.
- 3** Click **Choose selected** or drag and drop the selections in the drag area on the right side.
- 4** Click **Next** to submit your selections.
- 5** To remove a network pair, right-click the pair and select **Delete Pair**.

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

Managing Resiliency Platform Data Mover gateway pairing

This chapter includes the following topics:

- [Viewing Veritas Replication Gateways](#)
- [About Veritas Replication Gateway pairs](#)
- [How Resiliency Platform Data Mover supports encryption for data replication](#)
- [Creating a Veritas Replication Gateway pair](#)
- [Removing a Veritas Replication Gateway pair](#)
- [Viewing Veritas Replication Gateway pairs](#)
- [Modifying encryption for a Veritas Replication Gateway pair](#)

Viewing Veritas Replication Gateways

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

To view Replication Gateways

◆ Navigate



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

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

Click **Data Mover**

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

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

See [“About Veritas Replication Gateway pairs”](#) on page 61.

About Veritas Replication Gateway pairs

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

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

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

How Resiliency Platform Data Mover supports encryption for data replication

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

Creating a Veritas Replication Gateway pair

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

To create a Replication Gateway pair

1 Prerequisites

A Replication Gateway must be deployed in both data centers.

2 Navigate



Disaster Recovery Settings (navigation pane)

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

3 Specify the information in the wizard:

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

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

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

See [“Removing a Veritas Replication Gateway pair”](#) on page 62.

See [“About Veritas Replication Gateway pairs”](#) on page 61.

Removing a Veritas Replication Gateway pair

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

To remove a Replication Gateway pair

1 Prerequisites

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

2 Navigate



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

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

See [“About Veritas Replication Gateway pairs”](#) on page 61.

Viewing Veritas Replication Gateway pairs

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

To view Replication Gateway pairs

◆ Navigate



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

See [“About Veritas Replication Gateway pairs”](#) on page 61.

Modifying encryption for a Veritas Replication Gateway pair

When you create a Replication Gateway pair, you can specify an encryption scheme for replication. You can modify this option after a gateway pair is created.

To modify a Replication Gateway pair

1 Navigate



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

2 Select the vertical ellipses next to the pair name and select **Edit**.

3 In the wizard, change the encryption scheme selection and submit.

See [“About Veritas Replication Gateway pairs”](#) on page 61.

Configuring resiliency groups for remote recovery

This chapter includes the following topics:

- [Understanding the role of resiliency groups in disaster recovery operations](#)
- [How Resiliency Platform configures disaster recovery protection for virtual machines](#)
- [Managing virtual machines for remote recovery \(DR\)](#)
- [Verifying the replication status for Veritas Resiliency Platform Data Mover](#)
- [Deleting a resiliency group configured with Resiliency Platform Data Mover](#)

Understanding the role of resiliency groups in disaster recovery operations

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

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

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

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

See [“Managing virtual machines for remote recovery \(DR\)”](#) on page 67.

How Resiliency Platform configures disaster recovery protection for virtual machines

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

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

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

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

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

For VMware virtual machines, on successful completion of the operation, Resiliency Platform creates a directory in the working location of the virtual machine to save the virtual machine-related files for the recovery data center. Resiliency Platform uses these files during the DR operations such as Migrate, Takeover, Rehearsal, hence these files and the directory should not be deleted or modified. This directory lets you have separate configurations across the two data centers for the same virtual machines.

When you configure virtual machine for remote recovery using Resiliency Platform Data Mover, the existing storage policy is removed and again added back to the virtual machine. This may impact a few other rules.

Managing virtual machines for remote recovery (DR)

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

To manage virtual machines for remote recovery (DR)

1 Prerequisites

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

See [“DR for VMware virtual machines using Resiliency Platform Data Mover - an overview of key steps”](#) on page 37.

See [“DR for VMware virtual machines using array-based replication - an overview of key steps”](#) on page 39.

2 Navigate



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

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

- 3 Select the virtual machines:
 - Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.
 - Drag and drop virtual machines to **Selected Instances**.
 For VMware assets, if you plan to use Resiliency Platform Data Mover, each resiliency group must map to only one ESX cluster.
- 4 The next page displays the environment for the selected assets.
- 5 The next page lists the service objectives that are available for the selected asset type. You can expand the service objective to view details. Select the service objective that provides disaster recovery operations.

Note: If array-based replication is not configured, and if the selected asset type is supported by Resiliency Platform Data Mover, the wizard provides the option to use Resiliency Platform Data Mover.

- 6 Select the target (recovery) data center.
- 7 Continue through the wizard to complete the DR configuration. The information you must enter depends on the selected service objective and replication type:
 See [“DR configuration options for VMware virtual machines using Resiliency Platform Data Mover”](#) on page 68.
- 8 Complete the network customization steps for the virtualization technology.
 See [“Network customization options”](#) on page 70.
- 9 Verify the summarized information and enter a name for the resiliency group.
- 10 When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the virtual machines for the DR operation. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.
 See [“Viewing activities”](#) on page 98.
 Verify that the new resiliency group is added to the **Resiliency Groups** tab.
 See [“Viewing resiliency group details”](#) on page 24.

DR configuration options for VMware virtual machines using Resiliency Platform Data Mover

The following tables summarize the information that you must supply or verify in the wizard when you apply a service objective for a resiliency group.

Table 7-1 Options for Resiliency Platform Data Mover

Wizard steps and options	Description
Select Veritas Replication Gateway Pair	<p>Each gateway pairs provides information on the available assets and disks. If the gateways in the pair have different available capacities, the smaller capacity assets and disks are shown.</p> <p>Select a gateway pairs to use for this resiliency group.</p> <p>Ensure that the available virtual machine assets and disks are sufficient for the number of assets and disks listed as the Assets requirement.</p>
Select Target Cluster and Datastore	<p>Each resiliency group must map to only one ESX cluster. The wizard validates which clusters on the target data center can meet the required number of virtual machines and disks.</p> <p>Select the cluster in the Select Target ESX cluster.</p> <p>Select the datastore in the Select target Datastore.</p> <p>Once you select the cluster, the associated ESX hosts are displayed, and below them the datastores that are accessible from the ESX hosts and the selected gateway.</p> <p>Storage that can be provisioned must be available in the selected datastore in order for the wizard to create the replicated disks on the target data center. Review the GB of total data size in the Assets requirement list. Compare this to the value in the Free (GB) column for the selected datastore.</p> <p>If there is insufficient CPU or memory, you can continue with the wizard and update the resources later.</p>
Confirm Details	<p>Verify the gateway, cluster, and datastore selections before you continue with the wizard.</p>

Network customization options

You can do the following in this panel:

- Customize the static IP for virtual machines on the production and the recovery data center.
- Choose between Production and Rehearsal DNS customization.
- Manage PTR records
- Choose to continue with DR operations even if DNS updates fail.

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

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

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. This user name and password is to be specified under 'Windows global user'. The user credentials can be Windows Active Directory user or Workgroup user. For Windows Active Directory user, the Active Directory should be common for both, the primary and the recovery data center.

If a Windows virtual machine is part of a Windows Active Directory, ensure that you log on to the virtual machine at-least once using the Active Directory credentials.

For more information refer to the *Deployment Guide*.

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

See ["Prerequisites for network customization"](#) on page 70.

Prerequisites for network customization

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

- Linux and Windows virtual machines that are running on VMware platform ensure that the VMware Tools are installed and running.
- IP, Gateway, Netmask, DNS, Domain Name, Mac address etc. information should present in the respective files of each network interface for which you want to customize the IP and DNS.

- If multiple network interfaces (NICs) are assigned to a virtual machine, then you need to apply IP customization to all the NICs.
- For virtual machines that are running on Linux ensure that NetworkManager and libvirt service is in off state.
- The mac address configuration should be set as Manual/Static so that it does not change after the DR operation is performed.

Verifying the replication status for Veritas Resiliency Platform Data Mover

After configuring DR using Veritas Resiliency Platform Data Mover, Resiliency Platform starts replication for powered-on virtual machines. However, for powered-off virtual machines, first power-on these virtual machines to start replication. The replication process begins with a full synchronization of the data between the protected virtual machines on the source data center and the target data center. View the status of replication for a resiliency group from the resiliency group details page.

See [“Viewing resiliency group details”](#) on page 24.

To verify the replication status for Veritas Resiliency Platform Data Mover

1 Navigate



Assets (navigation panel)

Resiliency Groups

- 2 Double-click the resiliency group for which DR using Resiliency Platform Data Mover is already configured.
- 3 View the replication state at **Replication > State**.

The replication state is a combination of Data Mover state and Resiliency Platform replication state, the latter state in parenthesis - Data Mover state (Resiliency Platform replication state). The tables describes the function of each state.

Some possible states:

- Consistent | Active (Connected, Consistent)
- Consistent | Inactive (Connected, Inconsistent)
- Inconsistent | Not Syncing | Inactive (Disconnected, Inconsistent)
- Inconsistent | Syncing | Inactive (Connected, Consistent)

- Consistent | Stopped (Connected, Consistent)

Table 7-2 Data Mover replication states

Data Mover state	Description
Consistent or Inconsistent	Data state on the target data center.
Syncing or Not Syncing	The Syncing state represents that data is in inconsistent state and data transfer is in full synchronization mode.
Active or Inactive	Replication state on the target data center. Other possible replication states are: Stopped , Stopped on Target Forcefully , Aborted , or Frozen .

Table 7-3 Replication states

Resiliency Platform replication state	Description
Connected or Disconnected	Replication state of Resiliency Platform on the target data center.
Consistent, Inconsistent good, or Inconsistent Note that Inconsistent good state is not applicable to Data Mover.	Data state on the target data center.

Based on the Resiliency Platform replication states, note that some disaster recovery operations are restricted:

- Migrate operation: Is allowed when replication state is **Connected** and data state is **Consistent** or **Inconsistent Good**.
- Takeover operation: Is allowed when replication state is any and data state is **Consistent**.
- Rehearsal operation: Is allowed when replication state is **Connected** and data state is **Consistent**.
- Resync and Rehearsal Cleanup operations: Is allowed with all states.

Deleting a resiliency group configured with Resiliency Platform Data Mover

You can run this operation after successful completion of configure DR.

For invoking Resiliency Platform Data Mover, refer to the following points:

- During the operation, replication is stopped and replication sets are deleted on gateways and on source and target ESX hosts.
- All the virtual machines disks that are attached to the Replication Gateway are deleted.

To perform the delete operation on a configured resiliency group

- 1** Navigate



Assets (navigation panel)

Resiliency Groups

- 2** Double-click the resiliency group for which DR is already configured.
- 3** On the resiliency group DR Readiness page, click **Delete**.
- 4** On the **Delete**confirmation page, click **Submit**.

Rehearsing DR operations to ensure DR readiness

This chapter includes the following topics:

- [About ensuring the disaster recovery readiness of your assets](#)
- [Rehearse operations - array-based replication](#)
- [Rehearse operations - Resiliency Platform Data Mover](#)
- [Performing the rehearsal operation](#)
- [Performing cleanup rehearsal](#)

About ensuring the disaster recovery readiness of your assets

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

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

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

Rehearse operations - array-based replication

The requirements for rehearse operations for VMware virtual machines depend on the replication type.

[Rehearse operations with EMC SRDF-based replication](#)

[Rehearse operations with NetApp SnapMirror based replication](#)

Rehearse operations with EMC SRDF-based replication

- The device group should be associated with the snapshot LUNs. Resiliency Platform supports TimeFinder Snap and TimeFinder Mirror.
- Rehearsal operations for resiliency groups that are replicated using EMC SRDF technology in Asynchronous mode cannot be performed using TimeFinder Snap technology (VDEV devices). You need to configure TimeFinder Mirrors (BCV devices) to perform the rehearsal operations on such resiliency groups.
- When the rehearse operation is initiated, Resiliency Platform creates point-in-time snapshots, since rehearsal cannot work with existing snapshots.

Note: If there are any active snapshots that are in progress, you need to terminate the snapshots and refresh the asset discovery.

- The datastores on the snapshot device are attached on the DR host.
- Resiliency Platform registers the virtual machines in the production data center for rehearsal. They have identical configuration to the DR virtual machines, except these virtual machines consume storage from the datastore mounted using the snapshot volumes. These virtual machines are disconnected from the network and are unregistered during cleanup.

Rehearse operations with NetApp SnapMirror based replication

- NetApp SnapMirror based replication uses FlexClone for the rehearse operation, so the NetApp storage server must be enabled with the FlexClone license.
- When the rehearse operation is initiated, Resiliency Platform creates a point-in-time volume snapshot as part of the rehearsal operation. The snapshot volume is exported and mounted on the DR host.

Note: The rehearse operation breaks any ongoing replication between the source and destination storage server as the FlexClone operation cannot be performed on the destination read-only volume. SnapMirror replication resumes after the rehearsal cleanup operation.

- Resiliency Platform registers the virtual machines in the production data center for rehearsal. They have identical configuration to the DR virtual machines, except these virtual machines consume storage from the datastore that is mounted using the snapshot volumes. These virtual machines are disconnected from the network and are unregistered during cleanup.

Rehearse operations - Resiliency Platform Data Mover

Before you can perform a rehearse operation that uses Resiliency Platform Data Mover, map production network with rehearsal network. The configuration maps the source and rehearsal virtual machine. The rehearsal network simulates the production network environment so that the tests of the application and the workload on the rehearsal network represent a realistic scenario.

During the disaster recovery rehearsal, Veritas Resiliency Platform creates a temporary virtual machine in addition to the rehearsal virtual machine. The temporary virtual machine is created for storage optimization using VMWare's Linked Clone technology. Veritas Resiliency Platform creates a snapshot in the target data center of the virtual machines in the source data center. Resiliency Platform then provisions the virtual machines in the rehearsal network. The virtual machines are created from the snapshot. You can bring up the virtual machines in the rehearsal network to test that the failover works as expected. At any point in time, only one rehearse operation is allowed for a resiliency group.

When you are satisfied with the testing of the simulated failover to the target data center, use the rehearsal cleanup operation to clean up the rehearsal virtual machines in the resiliency group. The cleanup operation deletes all of the temporary objects that were created during the rehearsal. Without performing the rehearsal cleanup operation, the configuration does not allow to perform migrate or takeover operations.

Performing the rehearsal operation

Use the **Rehearsal** option on the Resiliency Platform console to perform the disaster recovery rehearsal, which verifies the ability of your configured resiliency group to

fail over to the disaster recovery (DR) data center during disaster. A rehearsal is a zero-downtime test that mimics the configuration, application data, storage, replication, and the fail over behavior of your resiliency group.

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

To perform the rehearsal operation

1 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 (VLANs) across the data centers is complete.
See [“Setting up network mapping between production and recovery data centers”](#) on page 58.
- Each type of replication has prerequisites and limitations for the rehearsal operation.
See [“Rehearse operations - array-based replication”](#) on page 75.
See [“Rehearse operations - Resiliency Platform Data Mover”](#) on page 76.

2 Navigate



Assets (navigation pane)

Resiliency Groups

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

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

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

See [“Performing cleanup rehearsal ”](#) on page 77.

Performing cleanup rehearsal

After you have performed the rehearsal operation successfully to verify the ability of your configured resiliency group to fail over on to the disaster recovery data center, you can use the cleanup rehearsal operation to clean up the rehearsal virtual

machines or applications in the resiliency group. All temporary objects created during the rehearsal operation are now deleted.

To perform cleanup rehearsal

1 Navigate



Assets (navigation pane)

Resiliency Groups

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

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

See [“Performing the rehearsal operation”](#) on page 76.

Performing disaster recovery operations

This chapter includes the following topics:

- [How Resiliency Platform Data Mover handles migrate](#)
- [Migrating a resiliency group of virtual machines](#)
- [How Resiliency Platform Data Mover handles takeover](#)
- [How Resiliency Platform Data Mover migrates workloads back to the source data center after a takeover](#)
- [Taking over a resiliency group of virtual machines](#)
- [Performing the resync operation](#)

How Resiliency Platform Data Mover handles migrate

Migration refers to a planned activity involving graceful shutdown of virtual machines at the source data center and bringing virtual machines up at the target data center and vice-versa. In this process, replication ensures that consistent virtual machine data is made available at the source and target data center.

When you migrate a resiliency group, the replication state should be connected and in a consistent state. When you initiate a Migrate operation, Veritas Resiliency Platform checks whether the virtual machines and source data center Replication Gateways are up. The migrate operation involves stopping virtual machines, detaching disks from virtual machines and attaching them to the gateway.

The migrate operation is similar to the takeover operation, except that the migrate operation is used when the virtual machines can be gracefully shut down. The takeover operation is used in case of disaster, when the source virtual machines are not reachable. The migrate operation involves shutting down the virtual machine from within the guest machine, whereas, the takeover operation involves only stopping the virtual machines.

There is difference between migrate and takeover operations. The migrate operation allows reverse replication, whereas the takeover operation does not allow reverse replication. After takeover, the resync operation is required to bring back data or workload onto the source data center.

Migrating a resiliency group of virtual machines

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

If a disaster requires a takeover by a recovery data center, you can later perform a migration from the recovery data center back to the production data center. In such a case, you must first use the resync operation to prepare for the migration.

See [“Performing the resync operation”](#) on page 82.

To migrate virtual machines

1 Prerequisites

For VMware virtual machines, ensure that the mapping of all the required port groups (VLANs) across the data centers is complete.

See [“Setting up network mapping between production and recovery data centers”](#) on page 58.

2 Navigate



Assets (navigation pane)

Resiliency Groups

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

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

How Resiliency Platform Data Mover handles takeover

Takeover is an activity initiated by a user when the source data center is down due to a disaster, and the virtual machines need to be brought up at the target (recovery) data center to provide business continuity.

The replication includes any changes to the boot disks of the virtual machines. As a result, bootable copies of the source virtual machines can be brought up on the recovery data center.

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. If the available data is within the acceptable limits, perform the takeover operation to bring up the source workloads in the recovery data center. The takeover operation provisions and brings up the virtual machines at the recovery data center using the latest data on the recovery data center storage.

The takeover activity operates on an entire resiliency group, even if the disaster affects only certain workloads in the resiliency group.

After a takeover, the virtual machine in the recovery data center runs the application and writes to the storage in the recovery data center.

How Resiliency Platform Data Mover migrates workloads back to the source data center after a takeover

A takeover operation as a result of a disaster event moves the workloads from a source (production) data center to a target (recovery) data center. You must plan for how to restore the workloads back to a production data center once it is up and running.

The first step in the process is to use the resync operation to ensure that the data in the target data center is synchronized to the source data center storage. The resync operation performs the full synchronization between the storage on the target and the source data centers. After the synchronization is complete, you can use the migrate operation to restore the virtual machines to the source data center.

Both the resync and migrate operations are performed for an entire resiliency group.

Taking over a resiliency group of virtual machines

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

To perform takeover operation on virtual machines

1 Prerequisites

For VMware virtual machines, ensure that the mapping of all the required port groups (VLANs) across the data centers is complete.

See [“Setting up network mapping between production and recovery data centers”](#) on page 58.

2 Navigate



Assets (navigation pane)

Resiliency Groups

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

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

Performing the resync operation

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

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

Use the Resync operation on the Resiliency Platform console to automate these steps for the required resiliency groups. This operation cleans up the residue which

includes stopping applications and virtual machines, unregistering virtual machines, unmounting file systems, datastores, etc.

Performing the resync operation

1 Prerequisites

Restart the ESX servers on primary site before performing resync operation. Restarting the ESX servers ensures that all stale references to virtual machines, disks, or datastores are released so that resync can work properly.

2 Navigate



Assets (navigation pane)

Resiliency Groups

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

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

Monitoring and reporting on assets status

This chapter includes the following topics:

- [About the Resiliency Platform Dashboard](#)
- [Understanding asset types](#)
- [Displaying an overview of your assets](#)
- [Viewing reports](#)

About the Resiliency Platform Dashboard

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

- Which of my data centers have Resiliency Platform managed assets?
- What is the mix of my assets by type and platform?
- Which assets are configured for disaster recovery?

The Dashboard has the following areas:

Global View

A world map that identifies the data centers that contain Resiliency Platform managed assets.

Lines between data centers indicate that replication takes place between the locations.

Mouse over an icon for basic Resiliency Platform configuration and asset configuration information for that data center. Click **More** for detailed information and recent activity.

Resiliency Groups and Virtual Business Services summaries

The upper right section of the dashboard displays total number of resiliency groups and virtual business services in the resiliency domain, as well as those at risk and normal.

Click a square in either the **Resiliency Groups** or **Virtual Business Services** summary to display a tab of detailed information.

Virtual Machines by Platform and OS

Displays a summary of virtual machines in all data centers or information on a single data center. Use the drop-down list to filter your results. The summary lists the virtual machine types by percentage and the platform types by number.

Application environment

Displays the number of applications and the application types. The chart shows the number of applications that are managed by InfoScale and those that are not managed by InfoScale.

Applications by Type

Displays a summary of application types in all data centers or in a single data center. Use the drop-down list to filter your results.

Top Resiliency Groups by Replication Lag

Ranks the resiliency groups according to how long it takes the recovery data center to be in sync with the active data center.

By Service Objective

Displays the percentage of virtual machines and applications that are unprotected or unmanaged.

Use the drop-down list to filter your results.

See [“Displaying resiliency group information and status”](#) on page 21.

Understanding asset types

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

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

Displaying an overview of your assets

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

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



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

- Unprotected resiliency groups, are groups under Resiliency Platform control, but that do not have disaster recovery configured.
See [“Managing virtual machines for basic monitoring”](#) on page 18.

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

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

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

For VBSs, the screen displays the following:

- The number of VBSs that are created from virtual machines and the number that are created from physical assets.
- The number of resiliency groups within the VBSs that are protected and the number that are only managed (not protected).

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>
Virtual Infrastructure Inventory	<p>Provides information about the virtual infrastructure inventory across data centers. A pie chart shows the platform and virtualization technology distribution of the virtual servers across all data centers.</p> <p>The Details table provides additional information for each virtual server.</p>
Activity Distribution History	Provides information about tasks, such as migrate, takeover, rehearse, start, and stop, performed for a specified duration.
Recovery Activity History by RG	Provides historical information about recovery tasks, such as migrate, takeover, rehearse, for each resiliency group.
Recovery Activity History by VBS	Provides historical information about recovery tasks, such as migrate, takeover, rehearse, for each VBS.

To view a report

1 Navigation

Click **Reports** (menu bar).

2 Do one of the following:

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

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

Monitoring risks

This chapter includes the following topics:

- [About risk insight](#)
- [Displaying risk information](#)
- [Predefined risks in Resiliency Platform](#)
- [Viewing the current risk report](#)
- [Viewing the historical risk report](#)

About risk insight

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

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

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

Risk insight generates two types of reports:

- **Current risk reports:** Provides the summary and detail information about all the current risks in your data center.
- **Historical risk reports:** Provides a summary and a detailed analysis of information about the risks in your environment during the specified period.

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

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

Table 11-1

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

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

Table 11-2

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

Displaying risk information

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

You can display risks in the following ways:

Table 11-3 Ways to display risks



To display ...	Do the following:
A complete list of risks across the resiliency domain	<ol style="list-style-type: none">On the menu bar, select  More Views > RisksOn the Risk page, double-click a risk in the table to display detailed information.

Table 11-3 Ways to display risks (*continued*)

To display ...	Do the following:
Risks that are associated with a specific resiliency group or virtual business service	<ol style="list-style-type: none">1 On the navigation pane, select  (Assets) and the tab for either Resiliency Groups or Virtual Business Services.2 On the tab, double-click a resiliency group or virtual business service to display detailed information.3 On the details page, note any risks that are listed in the At Risk area, and double-click the risk for details.

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

Predefined risks in Resiliency Platform

[Table 11-4](#) lists the predefined risks available in Resiliency Platform. These risks are reflected in the current risk report and the historical risk report.

Table 11-4 Predefined risks

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Veritas Infoscale Operations Manager disconnected	Checks for Veritas Infoscale Operations Manager to Resiliency Manager connection state	1 minute	Error	All operations	Check Veritas Infoscale Operations Manager reachability Try to reconnect Veritas Infoscale Operations Manager

Table 11-4 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
vCenter Password Incorrect	Checks if vCenter password is incorrect	5 minutes	Error	<ul style="list-style-type: none"> On primary site: start or stop operations On secondary site: migrate or takeover operations 	In case of a password change, resolve the password issue and refresh the vCenter configuration
VM tools not installed	Checks if VM Tools are not Installed. It may affect IP Customization and VM Shutdown.	Real time, when resiliency group is created	Error	<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 removed from Virtual Machine	Checks if snapshot has been removed from virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Edit the resiliency group to refresh configuration
Snapshot reverted on Virtual Machine	Checks if snapshot has been reverted on virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Remove and re-add the virtual machine to the Resiliency group by editing Resiliency group
Data Mover Daemon Crash	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX.	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Snapshot created on Virtual Machine	Checks if a snapshot has been created on Virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Edit the resiliency group to refresh configuration

Table 11-4 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
DataMover virtual machine in noop mode	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX.	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Resiliency group configuration drift	Checks if disk configuration of any of the assets in the resiliency group has changed.	30 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Resync 	Edit the resiliency group to first remove the impacted virtual machine from the resiliency group and then add it back to the resiliency group.
Global user deleted	Checks if there are no global users. In this case, the user will not be able to customize the IP for Windows machines in VMware environment.	Real time	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Takeover 	Edit the resiliency group or add a Global user
Missing heartbeat from Resiliency Manager	Checks for heartbeat failure from a Resiliency Manager.	5 minutes	Error	All	Fix the Resiliency Manager connectivity issue
Infrastructure Management Server disconnected	Check for Infrastructure Management Server(IMS) to Resiliency Manager(RM) connection state.	1 minute	Error	All	Check IMS reachability Try to reconnect IMS
Storage Discovery Host down	Checks if the discovery daemon is down on the storage discovery host	15 minutes	Error	Migrate	Resolve the discovery daemon issue

Table 11-4 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 ■ Takeover 	Edit the Resiliency Group and disable DNS customization
IOTap driver not configured	Checks if the IOTap driver is not configured	2 hours	Error	None	Configure the IOTap driver This risk is removed when the workload is configured for disaster recovery
VMware Discovery Host Down	Checks if the discovery daemon is down on the VMware Discovery Host	15 minutes	Error	Migrate	Resolve the discovery daemon issue
VM restart is pending	Checks if the VM has not been restarted after add host operation	2 hours	Error	Configure DR	Restart the VM after add host operation
New VM added to replication storage	Checks if a virtual machine that is added to a Veritas Replication Set on a primary site, is not a part of the resiliency group.	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Takeover ■ Rehearsal 	Add the virtual machine to the resiliency group.
Replication lag exceeding RPO	Checks if the replication lag exceeds the thresholds defined for the resiliency group. This risk affects the SLA for the services running on your production data center.	5 minutes	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Takeover 	Check if the replication lag exceeds the RPO that is defined in the Service Objective
Replication state broken/critical	Checks if the replication is not working or is in a critical condition for each resiliency group.	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Takeover 	Contact the enclosure vendor.

Table 11-4 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 ■ Takeover 	Unmount the mount point that is already mounted or is being used by other assets.
Disk utilization critical	Checks if at least 80% of the disk capacity is being utilized. The risk is generated for all the resiliency groups associated with that particular file system.	<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 ■ Takeover ■ Rehearsal 	Delete or move some files or uninstall some non-critical applications to free up some disk space.
ESX not reachable	Checks if the ESX server is in a disconnected state.	5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations 	Resolve the ESX server connection issue.
vCenter Server not reachable	Checks if the virtualization server is unreachable or if the password for the virtualization server has changed.	5 minutes	Error	<ul style="list-style-type: none"> ■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations 	<p>Resolve the virtualization server connection issue.</p> <p>In case of a password change, resolve the password issue.</p>

Table 11-4 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■ Takeover	Reduce the number of CPUs assigned to the virtual machines on the primary site to match the available CPU resources on failover target.
Host not added on recovery data center	Checks if the host is not added to the IMS on the recovery data center.	30 minutes	Error	Migrate	Check the following and fix: <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.

Viewing the current risk report

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

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

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

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

To view the current risk report

- 1 Navigation:
Click **Reports** (menu bar).
- 2 In the **Risk > Current Risk Report** section, click **Run** or **Schedule** to receive the report on the specified email address.

Viewing the historical risk report

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

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

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

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

- The various charts under **Risk by Category** display the open risks and new risks in the recoverability, continuity, and SLA categories at specific points of time within the duration specified by you.
- The **Resolving time chart** shows the average time to resolve the risk within the recoverability, continuity, and SLA categories.
- The **5 risks that took the longest time to resolve** chart shows the top 5 risks that took the longest time to be resolved, within the recoverability, continuity, and SLA categories. This information is displayed per resiliency group or per Virtual Business Service (VBS).

To view the historical risk report

- 1 Navigation:
Click **Reports** (menu bar).
- 2 In the **Risk > Risk History Report** section, click **Run** or **Schedule** to receive the report on the specified email address.

Managing activities and resiliency plans

This chapter includes the following topics:

- [Managing activities](#)
- [Managing resiliency plans](#)

Managing activities

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

See [“Viewing activities”](#) on page 98.

See [“Aborting a running activity”](#) on page 99.

Viewing activities

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

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

To view activities**1** Navigate**Activities** (menu bar).**2** Choose either of the following:

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

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

See [“Aborting a running activity”](#) on page 99.

Aborting a running activity

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

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

To abort an activity**1** Navigate

Do one of the following:

**Activities**. Skip to [2](#)

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

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

- Right click your mouse and click **Abort**.
- Click on the vertical ellipsis and select **Abort**

See [“Viewing activities”](#) on page 98.

Managing resiliency plans

Veritas Resiliency Platform provides a console for creating and customizing resiliency plans. The following topics cover how to create, edit, delete resiliency plan templates and resiliency plans and how to execute resiliency plans.

See [“About resiliency plans”](#) on page 100.

See [“Creating a new resiliency plan template”](#) on page 101.

See [“Editing a resiliency plan template”](#) on page 104.

See [“Deleting a resiliency plan template”](#) on page 105.

See [“Viewing a resiliency plan template”](#) on page 105.

See [“Creating a new resiliency plan”](#) on page 106.

See [“Editing a resiliency plan”](#) on page 107.

See [“Deleting a resiliency plan”](#) on page 108.

See [“Executing a resiliency plan”](#) on page 108.

See [“Viewing a resiliency plan”](#) on page 108.

See [“Creating a schedule for a resiliency plan”](#) on page 109.

See [“Editing a schedule for a resiliency plan”](#) on page 109.

See [“Deleting a schedule for a resiliency plan”](#) on page 110.

See [“Viewing a schedule for a resiliency plan”](#) on page 110.

About resiliency plans

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

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

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

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

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

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

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

See [“Creating a new resiliency plan template”](#) on page 101.

See [“Creating a new resiliency plan”](#) on page 106.

Creating a new resiliency plan template

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

- Start and stop a resiliency group.
- Rehearsal and cleanup rehearsal of a resiliency group.
- Migrate and takeover a resiliency group.
- Manual task
See [“About manual task”](#) on page 102.
- Run a custom script
See [“About custom script”](#) on page 103.

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

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

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

To create a new resiliency plan template

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Templates** section, click **New**.
- 3 In the **Create New Template** wizard panel, enter a name and a description for the template.
- 4 Drag and drop the required operation into the canvas. Connect the **Start** and **Stop** actions to the operation.
- 5 Click **Create**.

See [“About resiliency plans”](#) on page 100.

About manual task

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

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

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

Using manual tasks in resiliency plans

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

To use a manual task in a resiliency plan

- 1 You can add a manual task to a resiliency plan template or to a resiliency plan.
See [“Creating a new resiliency plan template”](#) on page 101.
See [“Creating a new resiliency plan”](#) on page 106.
- 2 Drag and drop **Manual Task** into the canvas. Click the pencil icon in the action box to add the task details.
- 3 Provide a name for the manual task.

- 4 Describe the reason why you want to add this manual task to the resilient plan.
- 5 Select your choice for resuming the process manually or automatically. If you select the option for automatically resuming the process after a timeout, enter the duration of timeout in minutes. Click **Save**.

About custom script

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

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

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

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

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

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

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

Using custom scripts in resiliency plans

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

To use a custom script execution task in a resiliency plan

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

See [“Creating a new resiliency plan template”](#) on page 101.

See [“Creating a new resiliency plan”](#) on page 106.

- 2 Drag and drop **Custom Script** into the canvas. Click the pencil icon in the action box to add the task details.
- 3 Enter a name for the custom script.
- 4 Select the data center and the host where you want to execute the script. Click **Next**.
- 5 Enter the following details:
 - The relative path of the script on the specified host. The script path that you enter is taken as relative to the `VRTSsfmh InstallDir/vrp/scripts/` directory path.
For example, if you enter the path of the script as `myscripts/backup_scripts/script_name`, then the complete path considered by the system will be `VRTSsfmh InstallDir/vrp/scripts/myscripts/backup_scripts/script_name`.
 - Command-line arguments to the script. This is an optional input field.
 - Timeout for the script. By default, there is no timeout for the script execution. You can specify a timeout for the script execution. After the specified timeout expires, the script execution task in the resiliency plan is marked as failure but the script execution task is not stopped. The script execution may continue in the background. If you do not specify any timeout, the task will wait till the script is not completed.
- 6 Click **Save**.

Editing a resiliency plan template

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

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

To edit a resiliency plan template

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
 - 2 In the **Templates** list, place your cursor on the row which you want to edit. Do one of the following:
 - Right click your mouse and click **Edit**.
 - Click on the vertical ellipsis and select **Edit**.
 - 3 In the **Edit Template** wizard panel, edit the required actions and click **Save**.
The steps for editing the plan are the same as creating it.
- See [“Creating a new resiliency plan template”](#) on page 101.

Deleting a resiliency plan template

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

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

To delete a resiliency plan template

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Templates** list, place your cursor on the row which you want to delete. Do one of the following:
 - Right click your mouse and click **Delete**.
 - Click on the vertical ellipsis and select **Delete**.
- 3 In the **Delete Template** panel click **Delete**.

See [“Creating a new resiliency plan template”](#) on page 101.

Viewing a resiliency plan template

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

To view a resiliency plan template

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Templates** list, do one of the following:
 - Double click the row that you want to view.
 - Select the row that you want to view, right click and select **Details**.
 - Select the row that you want to view, click on the vertical ellipsis and select **Details**.
- 3 You can now view the details of the resiliency plan template.

Creating a new resiliency plan

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

- Start and stop a resiliency group.
- Rehearsal and cleanup rehearsal of a resiliency group.
- Migrate and takeover a resiliency group.
- Manual task
See [“About manual task”](#) on page 102.
- Run a custom script
See [“About custom script”](#) on page 103.

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

To create a new resiliency plan using blank template

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Saved Plans** section, click **New**.
- 3 In the **Create Saved Plan - Select Template** wizard panel, select **Blank Template**, and click **Next**.
- 4 In the **Add Assets** panel, enter name and description.

- 5 Drag and drop the required operation into the canvas. Connect the **Start** and **Stop** actions to the operation.
- 6 Click the pencil icon in the action box to add relevant assets. Select the data center whose assets you want to add to the template. Click **Add**.
- 7 Click **Submit**.

To create a new resiliency plan using predefined template

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Saved Plans** section, click **New**.
- 3 In the **Create Saved Plan - "Select Template"** wizard panel, select **Pre-defined Template**.
- 4 Select a template from the list and click **Next**.
- 5 In the **Add Assets** panel, name and description are pre-populated.
- 6 Click the pencil icon in the action box to add relevant assets. Select the data center whose assets you want to add to the template. Click **Add**.
- 7 Click **Submit**.

See ["About resiliency plans"](#) on page 100.

See ["Deleting a resiliency plan"](#) on page 108.

See ["Executing a resiliency plan"](#) on page 108.

Editing a resiliency plan

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

To edit a resiliency plan

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Saved Plans** list, place your cursor on the row which you want to edit. Do one of the following:
 - Right click your mouse and click **Edit**.
 - Click on the vertical ellipsis and select **Edit**.
- 3 In the **Edit Saved Plan** wizard panel, edit the required actions and click **Submit**.
The steps for editing the plan are the same as creating it.

See ["Creating a new resiliency plan"](#) on page 106.

Deleting a resiliency plan

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

To delete a resiliency plan

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Saved Plans** list, place your cursor on the row which you want to delete. Do one of the following:
 - Right click your mouse and click **Delete**.
 - Click on the vertical ellipsis and select **Delete**.
- 3 In the **Delete Saved Plan** panel click **Delete**.

See [“Creating a new resiliency plan”](#) on page 106.

Executing a resiliency plan

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

To execute a resiliency plan

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Saved Plans** list, place your cursor on the row which you want to execute. Do one of the following:
 - Right click your mouse and click **Execute**.
 - Click on the vertical ellipsis and select **Execute**.
- 3 In the **Execute Saved Plan** panel click **Execute**.

See [“Creating a new resiliency plan”](#) on page 106.

Viewing a resiliency plan

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

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

See [“Editing a resiliency plan”](#) on page 107.

See [“Deleting a resiliency plan”](#) on page 108.

To view a resiliency plan

1 Navigate

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

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

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

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

Creating a schedule for a resiliency plan

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

To create a schedule for a resiliency plan

1 Navigate

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

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

- Double click the row for which you want to create a schedule. In the **Schedule** section of details page, click **New**.
- Select the row for which you want to create a schedule, right click and select **Create Schedule**.
- Select the row for which you want to create a schedule, click on the vertical ellipsis and select **Create Schedule**.

Editing a schedule for a resiliency plan

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

To edit a schedule for a resiliency plan

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Saved Plans** list, do one of the following:
 - Double click the row for which you want to edit a schedule. In the **Schedule** section of details page, click **Edit**.
 - Select the row for which you want to create a schedule, right click and select **Edit Schedule**.
 - Select the row for which you want to create a schedule, click on the vertical ellipsis and select **Edit Schedule**.

Deleting a schedule for a resiliency plan

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

To delete a schedule for a resiliency plan

- 1 Navigate
Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2 In the **Saved Plans** list, do one of the following:
 - Double click the row for which you want to delete a schedule. In the **Schedule** section of details page, click **Delete**.
 - Select the row for which you want to edit a schedule, right click and select **Delete Schedule**.
 - Select the row for which you want to edit a schedule, click on the vertical ellipsis and select **Delete Schedule**.

Viewing a schedule for a resiliency plan

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

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

See [“Editing a schedule for a resiliency plan”](#) on page 109.

See [“Deleting a schedule for a resiliency plan”](#) on page 110.

To view a schedule for a resiliency plan

- 1** Navigate
 Resiliency Plans (menu bar) or **Quick Actions > Resiliency Plans**
- 2** In the **Saved Plans** list, do one of the following:
 - Double click the row for which you want to view a schedule.
 - Select the row for which you want to view a schedule, right click and select **Details**.
 - Select the row for which you want to view a schedule, click on the vertical ellipsis and select **Details**.
- 3** In the **Schedule** section of details page, view the details of the schedule.

General troubleshooting

This appendix includes the following topics:

- [Viewing events and logs in the console](#)
- [Events in VMware virtual machines disaster discovery](#)
- [Troubleshooting discovery of assets](#)
- [Log files to troubleshoot Veritas Resiliency Platform Data Mover](#)
- [Resiliency Platform fails to attach storage policy to virtual machines](#)
- [Resiliency Platform fails to create storage policy](#)
- [Resolving the Admin Wait state](#)

Viewing events and logs in the console

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

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

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

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

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

To view events and logs

1 Navigate



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



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

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

Events in VMware virtual machines disaster discovery

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

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

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

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

Troubleshooting discovery of assets

When asset infrastructure is added to the Infrastructure Management Server (IMS), or when changes are made to the infrastructure, the IMS discovers and correlates

the asset information and displays the information on the Assets page of the Resiliency Platform console. The discovery can take some time before the information is updated on the console. Until discovery is complete, some information needed to configure resiliency groups may be missing from the Assets page on the console.

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

Note: Occasionally, the data discovered from the Infrastructure Management server (IMS) may not be updated properly in the Resiliency Manager database. This situation may result in displaying incorrect information about the resiliency group state, replication state, and replication type. In such a case, refresh the appropriate assets in both the data centers.

If you are configuring replication using storage arrays in a VMware vCenter Server environment, you can use the following guidelines to speed up discovery or to troubleshoot information that is not being updated:

Table A-1 Configuring asset infrastructure in IMS for storage arrays in VMware environment

Situation	Troubleshooting/best practices
Adding storage arrays as enclosures to IMS	Ensure that the storage arrays that are added to the IMS are the ones that provide storage to the ESX servers managed by the vCenter Server that is added to the IMS.
More than one IMS in data center	Ensure that the vCenter Server that is managing the ESX servers, and the enclosure providing storage to those servers, are added to the same IMS.
Refreshing the IMS after a change in infrastructure	Ensure that you use the Refresh operation on the correct vCenter Servers and enclosures where the change was made.
Refreshing the IMS after a change in infrastructure, where there is more than one IMS	Ensure that you use the Refresh operation in the correct IMS.

In the VMware and EMC SRDF environment, the general guideline is to add/refresh the enclosure before adding/refreshing the VMware vCenter Server.

Table A-2 Configuring or refreshing asset infrastructure in IMS for VMware and EMC SRDF environment

Situation	Recommended sequence
You have not yet added the asset infrastructure.	Add the enclosure information in the IMS and let the discovery complete before adding the vCenter Server to the IMS.
You later provision new storage from an enclosure that is already configured in the IMS and mount datastores from the new storage.	Refresh the enclosure in the IMS, let the refresh task on the enclosure complete, and then refresh the vCenter Server in the IMS.
You provision storage from a new enclosure.	Add the new enclosure in the IMS and then refresh the vCenter Server after the enclosure discovery completes.
You are provisioning storage from an enclosure that is already configured in the IMS to new ESX servers managed by a vCenter Server.	Refresh the enclosure first, then add the vCenter Server to the IMS or refresh it if it is already added to the IMS.

In the VMware and NetApp SnapMirror environment, the general guideline is add/refresh the vCenter Server first, then add/refresh the NetApp enclosure.

Table A-3 Configuring or refreshing asset infrastructure in IMS for storage arrays in VMware and NetApp SnapMirror environment

Situation	Recommended sequence
You have not yet added the asset infrastructure.	Add the vCenter Server to the IMS first and let the discovery complete before you add the NetApp enclosure.
You later provision storage from an existing NetApp enclosure and mount NFS datastores on ESX servers.	Refresh the vCenter Server first in the IMS, let the discovery complete and then refresh the NetApp enclosure.
You later provision storage from a new NetApp enclosure and mount NFS datastores on that ESX servers.	Refresh the vCenter Server first in the IMS, wait for the vCenter Server discovery to complete, and then add the new NetApp enclosure.

The recommended sequence for adding or modifying asset infrastructure for application discovery in the NetApp SnapMirror replication environment is as follows: Ensure that discovery of the hosts is complete before you add or refresh the NetApp enclosures.

For more information on adding asset infrastructure and on the refresh operation in the IMS, refer to the *Deployment Guide*.

Log files to troubleshoot Veritas Resiliency Platform Data Mover

Depending on the troubleshooting issue encountered, you can refer to particular log files to isolate and resolve issues.

Table A-4 Log files locations to troubleshoot Resiliency Platform Data Mover issues

Issue type	Log files	Location
VMware related operations, specifically, issues related to IO filter installation	/var/opt/VRTSsfmh/logs/vsphere_wsclient_op.log	IMS or Control Host
VMware storage policy related issues	/var/opt/VRTSsfmh/logs/vsphere_spbmclient_pm.log /var/opt/VRTSsfmh/logs/vsphere_spbmclient.log	IMS or Control Host
Operations performed on ESX hosts such as configuring Replication Set or start/stop replication and so on	/var/opt/VRTSsfmh/logs/vxvaio_controller.log	IMS or Control Host
Issues related to replication, start/stop operations	/var/log/iofilterd-vtstap.log	ESX host
Replication issues such as replication going into disconnected state and so on	<VM dir>/vmware.log For example, /vmfs/volumes/Shared/VM-1/vmware.log	vCenter
Storage policy related issues	VMware\vCenterServer\logs\vmware-sps\sps.log	vCenter
VMware operations, specifically, issues related to IO filter installation	VMware\vCenterServer\logs\vmware-vpx\vpzd-....log VMware\vCenterServer\logs\eam\eam.log	vCenter

Resiliency Platform fails to attach storage policy to virtual machines

The IO filter version added to the storage policy on the virtual machine and ESX host is different.

Error message: ESX host does not support the virtual hardware configuration of virtual machine. Check the `IMS/CH:`

`/var/opt/VRTSsfmh/logs/vsphere_wsclient_op.log` log file to verify and resolve the issue.

Resolution: Remove the storage policy from vSphere and perform the operation again.

Resiliency Platform fails to create storage policy

The clock time on IMS or Control Host and vSphere client is not synchronized.

Error message: Failed to perform operation. Log message in `vsphere_spbmclient_pm.log`: Client received SOAP Fault from server: The time now *Day Month Date Time Zone Year* does not fall in the request lifetime interval extended with clock tolerance of *milliseconds*. This might be due to a clock skew problem. See the server log to find more detail regarding exact cause of the failure.

Verify the following log files in IMS or Control Host server.

`/var/opt/VRTSsfmh/logs/vsphere_spbmclient.log`

`/var/opt/VRTSsfmh/logs/vsphere_spbmclient_pm.log`

Resolution: Ensure that the clock time on IMS or Control Host and vCenter server is synchronized and then re-run the workflow again.

Resolving the Admin Wait state

During replication of a resiliency group, an issue may occur which requires manual intervention for a Veritas Replication Set on the Replication Gateway. For example, there may be disk corruption or a disk may not be accessible. Any of the replication services, including the I/O Receiver, the Transceiver, the Applier, or the Replication Engine, may detect an issue when sending, receiving, or applying an update. The operation fails, and the replication state of the Resiliency Group shows in the web console as Admin Intervention Required.

To debug further, determine the type of issue and which Veritas Replication Set is affected by logging into the Replication Gateway and performing step 1 in the

procedure described below. A Veritas Replication Set consists of all the disks for a protected virtual machine.

After the issue is manually resolved, you must clear the Admin Wait state. On the Replication Gateway where the Veritas Replication Set has the Admin Wait state, use the following steps. This procedure assumes that the replication is from the production data center to the recovery data center. If the issue occurred during the prepare for failback operation, the replication direction is reversed.

To clear the Admin Wait state

- 1 You can clear the admin wait from recovery site Gateway. Login to the recovery site Gateway as admin user. This will invoke the CLISH interface:

```
monitor> datamover repl-sets
```

The output of this command shows the VRS-ID of the Veritas Replication Set. It also displays a column for Admin Intervention, which shows the Admin Wait string indicating the type of error and the service that reported the error.

See [“Admin Wait state codes ”](#) on page 118.

- 2 After resolving the underlying issue, clear the Admin Wait state on the Replication Set from recovery site Gateway using the following command of CLISH:

```
manage> datamover operation clear-admin-wait
```

You need to enter the VRS-ID of the Veritas Replication Set.

You also have an option to specify if you want a full sync to clear the Admin Wait.

- 3 Verify if the Admin Wait state has been resolved by using the following command of CLISH:

```
monitor> datamover repl-sets
```

Admin Wait state codes

The resiliency group shows the Admin Wait state if an issue occurs during replication of a protected virtual machine. The resiliency group requires manual intervention to fix the issue before the replication can resume.

Table A-5 Admin Wait state codes

Admin Wait string	Cause	Resolution	Needs full sync?
Engine: Local UpdateSetIDs misaligned Error	Internal errors in the replication solution.	Perform a full synchronization.	Yes.
Engine: Target Disk not found Error	After reboot, the target disk is not found.	Check the target disk and attach it. Clear the admin wait and resume replication.	No.
Transceiver: ConsistencyGroup is in invalid replication state on target gateway Error	Replication is not started on the target gateway.	Start the replication on the target gateway.	No.
Transceiver: Replication stopped on target gateway Error	Replication on the target gateway is stopped in the process of sending data from the source gateway (Internal error).	Start the replication on the target gateway.	No.
Transceiver: ConsistencyGroup not found on remote gateway Error	The consistency group is not configured on the remote gateway.	Delete the resiliency group and create a new resiliency group. Configure DR for the new resiliency group.	
Transceiver: Remote Gateway is not target gateway Error	After takeover, when the source gateway comes back online, both gateways can have the source role.	Run abort CG procedure for the on-premises Replication Gateway. Or run prepare for failback whenever it is ready to reverse sync from cloud Replication Gateway to the on-premises Replication Gateway.	
Transceiver: Disk Error	Opening and reading the update set files has failed.	Check <code>/var/log/messages</code> for any disk errors or file system errors. Check <code>vxgwaplrd</code> and <code>vxgwtxrxd</code> logs for any errors. Check whether the current update set is accessible. If the files cannot be restored even after manual intervention like <code>mount/fsck/check disk</code> , the disk must be replaced in the configuration and full synchronization is required.	Maybe.

Table A-5 Admin Wait state codes (*continued*)

Admin Wait string	Cause	Resolution	Needs full sync?
Transceiver: Data Corruption Error	Replication data files are corrupted	If the files are corrupted, resolution will require fsck/check disk. Perform a full synchronization.	Yes.
Applier: Disk Error	Opening and reading the update set files has failed.	Check <code>/var/log/messages</code> to see if there are any errors for disk or file system errors. Check <code>vxgwaplrd</code> and <code>vxgwtxrxrd</code> logs for any errors. Check whether the current update set is accessible. If the files cannot be restored even after manual intervention like mount/fsck/check disk, the disk must be replaced in the configuration and full synchronization is required.	Maybe.
Applier: Target Disk Error	Opening and writing on a target disk has failed	Check <code>/var/log/messages</code> to see if there are any errors for the target disk. If error cannot be resolved, the disk must be replaced in the configuration and full synchronization is required. If the issue occurred during prepare for failback, check the iSCSI connection, and check <code>/var/log/messages</code> for iSCSI errors. For network errors, full synchronization may not be needed. Login to Storage Proxy and check for <code>tgtd</code> status and errors.	Maybe.

Table A-5 Admin Wait state codes (*continued*)

Admin Wait string	Cause	Resolution	Needs full sync?
Applier: Target Disk not found Error	The target disk is not found on the gateway	<p>If the issue occurred during replication, check if the disk is attached to the cloud Replication Gateway. Also check OpenStack.</p> <p>If the issue occurred during prepare for failback, check that the disk is attached to the on-premises Replication Gateway. Also check iSCSI connection.</p> <p>Check <code>/var/log/messages</code> for any disk or virtio or virtblk or iSCSI errors.</p> <p>Check <code>vxgwaplrd</code> logs.</p> <p>If the disk was not attached, attach it and clear the admin wait flag.</p> <p>Full synchronization is not required if the update sets were not deleted. Otherwise, full synchronization is required.</p>	Maybe.
Applier: Data Corruption Error	Replication data files are corrupted.	<p>If the files are corrupted, the resolution requires <code>fsck</code> and check disk.</p> <p>Perform a full synchronization.</p>	Yes.
IOReceiver: Disk Error	Opening or reading or writing the update set files has failed.	<p>Check <code>/var/log/messages</code> for any disk errors or file system errors.</p> <p>Check <code>vxgwiorecvd</code> logs for any errors.</p> <p>Check whether the current update set is accessible.</p> <p>If the files cannot be restored even after manual intervention like <code>mount/fsck/check disk</code>, the disk must be replaced in the configuration and full synchronization is required.</p>	Maybe.
IOReceiver: Data Corruption Error	Replication data files are corrupted.	<p>If the files are corrupted, the resolution requires <code>fsck</code> and check disk.</p> <p>Perform a full synchronization.</p>	Yes.

Glossary

activity	A task or an operation performed on a resiliency group.
add-on	An additional software package that can be installed on hosts by the Infrastructure Management Server (IMS) for specialized uses.
asset infrastructure	The data center assets that can be added to the Infrastructure Management Server (IMS) for IMS discovery and monitoring. For example, virtual machines or virtualization servers.
assets	In Veritas Resiliency Platform, the virtual machines or applications that have been discovered by the Infrastructure Management Server (IMS) and that can be grouped into resiliency groups.
CLISH	Command Line Interface SHell. Provides the command line menu on the virtual appliance for use after the initial bootstrap configuration.
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 production data center and recovery data center. Each data center has a Resiliency Manager and one or more IMSs.</p>
host	<p>Physical servers, virtual machines, or Hyper-V servers that are added to the Infrastructure Management Server (IMS) as hosts.</p> <p>Adding the assets as hosts installs the host package that is used by the IMS for discovery and monitoring.</p>
Infrastructure Management Server (IMS)	The Veritas Resiliency Platform component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager.
migrate	A planned activity involving graceful shutdown of virtual machines at the production data center and starting them at the recovery data center. In this process, replication ensures that consistent virtual machine data is made available at the recovery data center.
persona	A user role that has access to a predefined set of jobs (operations). Used to assign permissions to users and groups for Veritas Resiliency Platform web console operations.
product role	The function configured for a Veritas Resiliency Platform virtual appliance.

	For example, a virtual appliance can be configured as a Resiliency Manager, Infrastructure Management Server (IMS) or both.
production data center	The data center that is normally used for business. See also recovery data center.
recovery data center	The data center that is used if a disaster scenario occurs. See also production data center.
rehearsal	<p>A zero-downtime test that mimics the configuration, application data, storage, and the failover behavior of the resiliency group.</p> <p>Rehearsal verifies the ability of the resiliency group to fail over to the recovery data center during a disaster.</p>
resiliency domain	The logical scope of a Resiliency Platform deployment. It can extend across multiple data centers.
resiliency group	The unit of management and control in Veritas Resiliency Platform. Related assets are organized into a resiliency group and managed and monitored as a single entity.
Resiliency Manager	The Veritas Resiliency Platform component that provides resiliency capabilities within a resiliency domain. It is composed of loosely coupled services, a distributed data repository, and a management console.
resiliency plan	A collection of tasks or operations, along with the relevant assets, which are performed in a predefined sequence.
resiliency plan template	A template defining the execution sequence of a collection of tasks or operations.
take over	An activity initiated by a user when the production data center is down due to a disaster and the virtual machines need to be restored at the recovery data center to provide business continuity.
tier	<p>Within a virtual business service (VBS), resiliency groups are arranged as tiers. Tiers represent the logical dependencies between the resiliency groups and determine the relative order in which the resiliency groups start and stop.</p>
virtual appliance	<p>An appliance that includes the operating system environment and the software application which are deployed together as a virtual machine.</p> <p>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).</p>
virtual business service (VBS)	A multi-tier IT service where each VBS tier hosts one or more resiliency groups. A VBS groups multiple services as a single unit for visualization, automation, and controlled start and stop in the desired order. You can also migrate/takeover the entire VBS.
web console	The web-based management console on the Resiliency Manager that is used to configure the settings for the resiliency domain and perform operations.

Index

A

- access profile
 - configuring for a data center 55
- activities
 - abort 99
 - view 98
- asset infrastructure
 - troubleshooting discovery of assets 113
- asset types
 - about 86
- assets
 - configuring for monitoring 18
 - display overview 86

B

- Bind settings for data center 57

D

- dashboard 84
- Data Mover
 - key configuration steps 37
- data replication
 - after takeover 81
 - handling application writes 33
 - handling migration 79
 - handling rehearsal 76
 - handling takeover 81
 - synchronizing data 32
- disaster recovery
 - configuration 39
 - limitations 54
 - prerequisites for virtual machines 50
 - Resiliency Platform Data Mover 37
 - using Resiliency Platform 10
- disaster recovery operations
 - cleanup rehearsal 77
 - migrate 80
 - rehearse 76
 - rehearse operations 74–75
 - resync 82

- disaster recovery operations *(continued)*
 - takeover 82
- DNS server settings for data center 57

E

- encryption for data replication 61
- events 112–113

F

- full synchronization
 - about 32

I

- Infrastructure Management Server
 - troubleshooting discovery of assets 113

L

- logs
 - viewing in console 112

M

- migration
 - about 79
- monitoring
 - events 113

N

- network customization
 - options 70
 - prerequisites 70
- network settings
 - configuring for a data center 55
 - editing for a data center 56

P

- permissions
 - about 14
- physical machines
 - configuring for remote recovery 67

R

- RBT disk
 - see Replication Block Tracking disk 33
- rehearse operations 74–77
- replication
 - EMC RecoverPoint 45
 - EMC SRDF 41
 - Hitachi True copy 47
 - monitoring events 113
 - NetApp SnapMirror 43
- Replication Block Tracking disk
 - about 33
- Replication Gateway
 - about 61
 - creating pair 61
 - modifying encryption scheme 63
 - removing pair 62
 - viewing 60
 - viewing pair 63
- replication lag 24
- replication states 71
- replication status 24
- reports
 - current risk 96
 - historical risk 97
 - viewing 87
- resiliency groups
 - about 16
 - deleting 25
 - displaying information and status 21
 - editing 24
 - guidelines for organizing 17
 - roles 65
 - starting 19
 - stopping 20
 - viewing detailed information 24
- resiliency plan template
 - viewing 105
- resiliency plan templates
 - create 101
 - deleting 105
 - editing 104
- resiliency plans
 - about 100
 - create schedule 109
 - creating 106
 - custom script 103
 - delete schedule 110
 - deleting 108

- resiliency plans *(continued)*
 - edit schedule 109
 - editing 107
 - executing 108
 - manual task 102
 - view schedule 110
 - viewing 108
- Resiliency Platform
 - features and components 11
- Resiliency platform
 - capabilities 13
- risk information
 - view 90
- risk insight
 - about 89
- risks
 - description 91

S

- service objectives
 - about 17
- states
 - replication 71
- subnet mapping 58
- synchronizing data
 - about 32

T

- takeover
 - about 81
- Troubleshoot Veritas Resiliency Platform Data Mover
 - Log files 116

V

- virtual machines
 - configuring for monitoring 18
 - configuring for remote recovery 67
- VMware vCenter Server privileges 52
- vtstap
 - about 33