

Veritas™ Resiliency Platform 1.2: Solutions for VMware

Veritas Resiliency Platform: Solutions for VMware

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

Product version: 1.2

Document version: 1.2 Rev 0

Legal Notice

Copyright © 2016 Veritas Technologies LLC. All rights reserved.

Veritas and the Veritas Logo 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. Please see the Third Party Legal Notice Appendix to this Documentation or TPIP ReadMe File accompanying this product for more information on the Third Party Programs.

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. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within the company to answer your questions in a timely fashion.

Our support offerings include the following:

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

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

www.veritas.com/support

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

Contacting Technical Support

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

www.veritas.com/support

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

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

- Product release level
- Hardware information
- Available memory, disk space, and NIC information

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

Licensing and registration

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

www.veritas.com/support

Customer service

Customer service information is available at the following URL:

www.veritas.com/support

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

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

Support agreement resources

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Contents

Technical Support	4
Chapter 1	Overview of Resiliency Platform 11
	About Veritas Resiliency Platform 11
	About Resiliency Platform features and components 12
	Resiliency Platform capabilities 13
	About permissions for operations in the console 14
Chapter 2	Managing VMware virtual machines using Resiliency Platform 15
	About managing VMware virtual machines using Resiliency Platform 15
Chapter 3	Managing resiliency groups 17
	About resiliency groups 17
	Guidelines for creating resiliency groups 18
	Managing and monitoring virtual machines 18
	Protecting virtual machines 20
	Displaying resiliency group information and status 23
	Displaying resiliency group details 25
	Modifying a resiliency group 27
	Starting a resiliency group 28
	Stopping a resiliency group 29
	Deleting a resiliency group 30
Chapter 4	Monitoring and reporting on assets status 31
	About the Resiliency Platform Dashboard 31
	Understanding asset types 33
	Displaying an overview of your assets 33
	Viewing reports 34

Chapter 5	Using Resiliency Platform for disaster recovery	36
	About disaster recovery using Resiliency Platform	36
	Understanding the role of resiliency groups in disaster recovery operations	37
Chapter 6	VMware virtual machines disaster recovery	38
	Understanding virtual machine disaster recovery	38
	Limitations for virtual machine disaster recovery	38
	VMware virtual machines disaster recovery - an overview of key steps	40
Chapter 7	Preparing for disaster recovery operations	42
	Configuring subnet information for a data center	43
	Setting up subnet mapping between production and recovery data centers	43
	Setting up virtual switch mapping between production and recovery data centers	44
	Configuring disaster recovery for a resiliency group of virtual machines	45
	Configuring disaster recovery - DR data center selection	47
	Configuring disaster recovery - Summary	47
	Viewing the details of a disaster recovery-enabled resiliency group	48
Chapter 8	Rehearsing DR operations to ensure DR readiness	49
	Ensuring the disaster recovery readiness of your Resiliency Platform assets using the rehearse operation	49
	Rehearse operations in VMware virtual machines	50
	Performing rehearsal cleanup	51
Chapter 9	Monitoring risks	53
	About risk insight	53
	Setting up replication lag threshold	54
	Displaying risk information	55
	Predefined risks in Resiliency Platform	55
	Viewing the current risk report	58
	Viewing the historical risk report	59

Chapter 10	Performing disaster recovery operations	60
	Migrating a resiliency group of virtual machines	60
	Taking over a resiliency group of virtual machines	61
	Performing the resync operation	61
Chapter 11	Managing activities and resiliency plans	63
	Managing activities	63
	Viewing activities	63
	Aborting a running activity	64
	Managing resiliency plans	65
	About resiliency plans	65
	Creating a new resiliency plan template	66
	Editing a resiliency plan template	70
	Deleting a resiliency plan template	70
	Viewing a resiliency plan template	71
	Creating a new resiliency plan	71
	Editing a resiliency plan	72
	Deleting a resiliency plan	73
	Executing a resiliency plan	73
	Viewing a resiliency plan	74
	Creating a schedule for a resiliency plan	74
	Editing a schedule for a resiliency plan	75
	Deleting a schedule for a resiliency plan	75
	Viewing a schedule for a resiliency plan	76
Appendix A	Configuring VMware environment for disaster recovery using replication	77
	Configuring VMware virtual machines disaster recovery using EMC SRDF replication	77
	Configuring VMware disaster recovery using NetApp SnapMirror	80
	Configuring VMware virtual machines disaster recovery using EMC RecoverPoint replication	83
	Configuring VMware virtual machines disaster recovery using Hitachi TrueCopy/ Hitachi Universal Replicator replication	86
	Configuring VMware virtual machines disaster recovery using HPE 3PAR Remote Copy replication	88
Appendix B	Troubleshooting	92
	Troubleshooting discovery of assets	92
	Viewing events and logs in the console	94

Events in VMware virtual machines disaster discovery	95
Glossary	97
Index	99

Overview of Resiliency Platform

This chapter includes the following topics:

- [About Veritas Resiliency Platform](#)
- [About Resiliency Platform features and components](#)
- [Resiliency Platform capabilities](#)
- [About permissions for operations in the console](#)

About Veritas Resiliency Platform

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

Resiliency Platform has the following core capabilities:

Recovery	Resiliency Platform provides a disaster recovery (DR) solution using data centers on premises in different geographical locations. The management console simplifies recovery, with single-click rehearsal and recovery operations.
Visibility	The console Dashboard provides visibility into the health of applications, virtual machines, and multi-tier business services.
Orchestration	Resiliency Platform can assist in data center day-to-day workload automation activities. For instance, virtual machines or IT services can be started and stopped for maintenance.

About Resiliency Platform features and components

The following is a brief introduction to Veritas Resiliency Platform key components and features. Administrators responsible for deploying and configuring 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>
data center	<p>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.</p>
asset infrastructure	<p>The data center assets that you add to the IMS for discovery and monitoring.</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	<p>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.</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 migrate or takeover the entire VBS.

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
Protecting and managing virtual machines as a single entity.	See “Managing and monitoring virtual machines” on page 18.
Displaying an overview of your resiliency domain including the number and health of your resiliency groups.	See “About the Resiliency Platform Dashboard” on page 31. See “Displaying resiliency group information and status” on page 23.
Starting and stopping resiliency groups for maintenance.	See “Starting a resiliency group” on page 28. See “Stopping a resiliency group” on page 29.
Configuring disaster recovery for a resiliency group	See “Configuring disaster recovery for a resiliency group of virtual machines” on page 45.
Monitoring risks for protected assets	See “About risk insight” on page 53.
Rehearsing disaster recovery	See “Ensuring the disaster recovery readiness of your Resiliency Platform assets using the rehearse operation” on page 49.
Migrating a resiliency group	See “Migrating a resiliency group of virtual machines” on page 60.
Taking over resiliency groups	See “Taking over a resiliency group of virtual machines” on page 61.
Viewing reports	See “Viewing reports” on page 34.

Table 1-1 Resiliency Platform capabilities (continued)

Capability	More information
Managing activities and resiliency plans	See “Managing activities” on page 63. See “Managing resiliency plans” on page 65.

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 VMware virtual machines using Resiliency Platform

This chapter includes the following topics:

- [About managing VMware virtual machines using Resiliency Platform](#)

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.

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

See [“About resiliency groups”](#) on page 17.

Note: Make sure you install VMware tools for VMware virtual machines.

Using Resiliency Platform capabilities to perform workload management tasks on virtual machines

The Resiliency Platform capabilities allow you to perform the tasks required for routine maintenance activities. For example, stop a resiliency group that contains a set of related virtual machines, update the required software components, and then restart the resiliency group.

Using Resiliency Platform to protect your VMware virtual machines

Resiliency Platform provides disaster recovery operations for VMware virtual machines. For example, migrate your resiliency group to another data center or perform the rehearse (fire-drill) operation to ensure that your infrastructure is adequately prepared for protection in the event of disaster. The replication required for these operations is provided by storage array.

See [“VMware virtual machines disaster recovery - an overview of key steps”](#) on page 40.

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

Managing resiliency groups

This chapter includes the following topics:

- [About resiliency groups](#)
- [Managing and monitoring virtual machines](#)
- [Protecting virtual machines](#)
- [Displaying resiliency group information and status](#)
- [Displaying resiliency group details](#)
- [Modifying a resiliency group](#)
- [Starting a resiliency group](#)
- [Stopping a resiliency group](#)
- [Deleting a resiliency group](#)

About resiliency groups

In Veritas Resiliency Platform, you organize related assets into a resiliency group that you can protect and manage as a single entity. Before you create a resiliency group, you must add the assets to Resiliency Platform.

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.

You can create a resiliency group in the following ways:

- You can create a resiliency group without enabling disaster recovery for it. See [“Managing and monitoring virtual machines”](#) on page 18.
- You can create a resiliency group and enable disaster recovery for it. This is known as a protected resiliency group. See [“Protecting virtual machines”](#) on page 20.

Guidelines for creating resiliency groups

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

If you create a resiliency group of virtual machines, follow these guidelines for selecting virtual machines:

- They reside on the same virtualization server or virtualization servers within the same VMware cluster.
They reside on the same ESX server.
- They all consume storage from the same replication consistency group (Symmetrix Remote Data Facility (SRDF) device group or NetApp volume).

Managing and monitoring virtual machines

A resiliency group lets you manage and monitor a group of assets as a single entity. For example, when you start a resiliency group, you start all the assets in the group.

You can create a resiliency group from virtual machines or applications, but not a mix of both.

You can organize any group of virtual machines into a resiliency group. However, the virtual machines often have a common characteristic. For example, they are all members of the same consistency group or they have the same virtualization server or hypervisor.

To create a resiliency group from virtual machines

1 Prerequisites

The asset infrastructure for the virtual machines must be added to the Infrastructure Management Server (IMS) and IMS discovery must be complete.

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

2 Navigate



Assets > Resiliency Groups tab > Manage & Monitor Virtual Machines or Applications

Or

Assets > Unmanaged tab > Manage & Monitor Virtual Machines or Applications

3 Display a list of virtual machines

On the **Select Assets** screen, use one or more of the following drop-downs to filter your list of virtual machines:

Asset Type	Select Virtual Machine .
Data Center	The data center in which the virtual machine is located.
Virtualization	The virtualization type.

4 Filter the list of virtual machines (optional)

Group By	Organize the virtual machines by virtualization server or replication consistency group.
Search	If you have a long list of virtual machines, use the Search field to filter the list.
show assets in resiliency group	When you select this check box, the list of virtual machines is updated with a Resiliency Group column. If a virtual machine is already a member of a resiliency group, this column displays the name of the group.

5 Select the virtual machines

To include a virtual machine in your new resiliency group, drag it from the list and drop in the **Selected Instances** area. To unselect a virtual machine from the **Selected Instances** area, you can drag it back to the list of virtual machines. When you select all the assets you need, click **Next**.

6 Create the resiliency group

Review the list of virtual machines that form your new resiliency group. If you need to make any changes, click **Back** to return to the **Select Assets** screen. When you are ready, name the resiliency group and click **Submit**.

7 Verify

On the confirmation screen, click **Done**.

A screen is displayed showing detailed information about the new resiliency group. It includes the following:

- The active data centers, replication type, and replication state.
- Controls to modify, delete, start, and stop the resiliency group.
- The disaster recovery readiness of the resiliency group. You can configure disaster recovery from this screen.
- A list of the virtual machines in the resiliency group.
- A list of risks (if any) to the resiliency group.

Click **Recent Activities** (bottom pane), and click **Details** to view the details of this task in a graphical representation.

More information is available on troubleshooting discovery of virtual machines.

See [“Troubleshooting discovery of assets”](#) on page 92.

Protecting virtual machines

Veritas Resiliency Platform lets you protect your virtual machines by creating a resiliency group and setting up disaster recovery for the group in a single set of steps.

Note: Even if you create a resiliency group without disaster recovery (using the Manage & Monitor Virtual Machines option), you can still configure disaster recovery later. On the **Assets** page, **Resiliency Groups** tab, right click the resiliency group and select **Configure DR**.

See [“Configuring disaster recovery for a resiliency group of virtual machines”](#) on page 45.

To protect virtual machines

1 Prerequisites

- The virtual machines you use to create the resiliency group must reside in two data centers: the production data center and the recovery data center. The asset infrastructure for the virtual machines must be added to the Infrastructure Management Server (IMS) at the appropriate data center and IMS discovery of the virtual machines must be complete.
For more information on adding asset infrastructure, refer to the *Deployment Guide*.
- If a VMware virtual machine has more than one ethernet adapters then all of them should have either static IP configuration or DHCP IP configuration. Mix of static and DHCP IP configuration is not supported on the same virtual machine.
- For a VMware 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.

2 Navigate



Assets > Resiliency Groups tab > Protect Virtual Machines

You can also access the **Protect Virtual Machines** wizard from the **Quick Action** menu.

3 Display a list of virtual machines

On the **Select Assets** screen, use one or more of the following drop-downs to filter your list of virtual machines:

Data Center	The data center in which the virtual machine is located.
Virtualization	The virtualization type.

4 Filter the list of virtual machines (optional)

Group By	Organize the virtual machines by virtualization server or replication consistency group.
Search	If you have a long list of virtual machines, use the Search field to filter the list.
show assets in resiliency group	When you select this check box, the list of virtual machines is updated with a Resiliency Group column. If a virtual machine is already a member of a resiliency group, this column displays the name of the group.

5 Select the virtual machines

To include a virtual machine in your new resiliency group, drag it from the list and drop in the **Selected Instances** area. If you change your mind, you can drag it back to the list of virtual machines. When you select all the assets you need, click **Next**.

6 Under **Manage Assets** screen, create the resiliency group.

Review the list of virtual machines that form your new resiliency group. If you need to make any changes, click **Back** return to the **Select Assets** screen. When you are ready, name the resiliency group and click **Next**.

7 Under **Configure DR** screen, configure DR for the selected resiliency group.

When you have configured DR for the resiliency group, click **Next**.

8 Select your disaster recovery data center

The **Select DR Datacenter** screen identifies your active data center, and lists the data centers you can select for disaster recovery. When you select a disaster recovery data center for your resiliency group, make sure that the data center has copies of the same virtual machines. Select the check box for the data center you want use and click **Next**.

9 Confirm that there are virtual machines at the recovery data center that match the virtual machines in your resiliency group.

In the **VM Selection** screen, verify the VM configurations for the production data center and recovery data center.

10 Complete the configuration

The **Summary** screen lists the following:

- The data center in which the resiliency group is located
- The recovery data center you specified

- The number of virtual machines in the recovery data center that you need to enable

Optionally, you can use the **Summary** screen to apply customized network settings, provided the subnets across the data centers are mapped.

If the configuration information is accurate and complete, click **Submit**.

11 Verify

On the confirmation screen, click **Done**.

The **Resiliency Group** tab is displayed, showing the new resiliency group.

You can use the **Quick Actions** drop-down list to perform other Veritas Resiliency Platform tasks.

For VMware virtual machines, on successful completion of the operation, the 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.

See [“Troubleshooting discovery of assets”](#) on page 92.


Displaying resiliency group information and status

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

Table 3-1 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 31.

Table 3-1 Displaying resiliency group information and status *(continued)*

Location	Level of detail	Useful for
 Assets > Resiliency Groups tab	Medium. Lists all your resiliency groups in one place.	Seeing what is in each of your data centers, the state of the groups, whether disaster recovery is configured, and so on.
Resiliency group-specific screen	Highest. Lists each asset in the resiliency group, their type, and state.	Getting detailed information on a resiliency group and its underlying assets. This screen can help you decide whether to start, stop, edit, or delete a group. See "Displaying resiliency group details" on page 25.

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 > Resiliency Groups tab

2 Review information and status

- For a quick health check of your resiliency groups, review the colored boxes above the table. Click on a box to show only the resiliency groups in that category; for example, click 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 click on a table heading to sort the groups. In the table, the key fields are **State**, **DR Status**, and **Replication Type**. Possible states are:

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>
DR Status	<p>Configured - The resiliency group has been configured for disaster recovery.</p> <p>Not Configured - Disaster recovery is not configured for the group. Configure it as soon as possible.</p>
Replication Type	<p>Resiliency Platform supports several replication technologies.</p> <p>If no replication type is shown, consider configuring replication.</p>
Type	<p>Application Group: The resiliency group comprises applications.</p> <p>Virtual Machine Group: The resiliency group comprises virtual machines.</p>

3 Display detailed information on a resiliency group (optional)

To display detailed information about a resiliency group, click its row in the table.

See [“Displaying resiliency group details”](#) on page 25.

Displaying resiliency group details

You can display detailed information on each of your resiliency groups. You can use a resiliency group-specific screen to answer questions as such the following:

- What is the overall health of the resiliency group?
- Is it configured for disaster recovery (DR)?
- What are its underlying assets and their current state?
- If DR is configured for the resiliency group, what is the replication lag time between sites?

To display details on a single resiliency group

1 Navigate



Assets > Resiliency Groups tab

2 Sort and select your resiliency group

On the **Resiliency Groups** tab, use the drop-down list, **Search** field, and table headings to filter your list of resiliency groups.

3 Display the resiliency group-specific screen

Double-click the table row for the resiliency group you are interested in.

The screen is divided into the following areas:

Table 3-2 Resiliency group details screen

This part of the screen ...	Displays ...
Top	<p>Resiliency group's health and status.</p> <p>It identifies the data centers at which the resiliency group is active, its replication state and type, and whether the resiliency group is configured for disaster recovery. This part of the screen displays the number of alerts that are associated with the resiliency group. And also verifies whether the resiliency group is part of any VBS or not.</p>
Middle	<p>A table with the assets that make up the resiliency group. You can use links above the table to sort the assets by data center, and you can use the table headings to sort the assets by Name, Type, or State.</p>
Bottom	<p>If the resiliency group is configured for disaster recovery, this portion of the screen displays the replication lag between the production data center and the recovery data center, and the recovery time. Note that the recovery time is available only after the rehearse operation is complete.</p>

You also can display information on your resiliency groups in the following ways:

- For a high-level view of resiliency group health, use the Resiliency Platform Dashboard.
See [“About the Resiliency Platform Dashboard”](#) on page 31.
- For a list of your resiliency groups and a quick view of which ones are up, configured, and so on, use the **Assets > Resiliency Group** tab.
See [“Displaying resiliency group information and status”](#) on page 23.

Modifying a resiliency group

You can modify resiliency group information including the group name as well as change the underlying assets on which the resiliency group is based.

Note: If you modify a resiliency group that has been configured for disaster recovery, you must reconfigure it.

To modify resiliency group information

1 Prerequisites

- After you configure a resiliency group for disaster recovery, you cannot edit the resiliency group. You must first unconfigure disaster recovery for the resiliency group, edit it, and then configure disaster recovery again.
- Determine the potential impact modifying the resiliency group may have on users. If necessary, notify users of the upcoming change.

2 Navigate



Assets > Resiliency Groups tab

3 Select

Use the on-screen filters, **Search** field, and table heading sort feature to locate your resiliency group.

4 Edit

Do one of the following:

- Right click on the resiliency group row and select **Modify**.
- On the right side of the resiliency group row, click on the vertical ellipsis and select **Modify**.

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

When you edit a resiliency group made up of virtual machines, note the following:

- If the resiliency group is configured for disaster recovery, Resiliency Platform proceeds to the Protect VM wizard.
- When the number of virtual machines on the replicated volume changes, edit the resiliency group to add or remove the virtual machines.

Note: If you add or remove virtual machines from a resiliency group after you have configured DR for that particular resiliency group, the DR functionality may not work as expected. You need to reconfigure DR for the resiliency group with the current set of virtual machines.

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

Starting a resiliency group

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

To start a resiliency group

1 Prerequisites

Create a resiliency group.

2 Navigate



Assets > Resiliency Groups tab

3 Select

Use the on-screen filters, **Search** bar, and table heading sort feature to locate your resiliency group.

4 Start the resiliency group.

Do one of the following:

- Right click on the resiliency group row and select **Start**.
- On the right side of the resiliency group row, click on the vertical ellipsis and select **Start**.

5 On the **Start Resiliency Group** screen, select the data center in which to start the group and click **Submit**.

6 Confirm

Click **Done**.

7 Notify

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

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

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 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 affect 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 stop the resiliency group.

2 Navigate



Assets > Resiliency Groups tab

3 Select

Use the on-screen filters, **Search** field, and table heading sort feature to locate your resiliency group.

4 Stop the resiliency group.

Do one of the following:

- Right click on the resiliency group row and select **Stop**.
- On the right side of the resiliency group row, click on the vertical ellipsis and select **Stop**

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

6 Confirm

Click **Done**.

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

Deleting a resiliency group

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

To delete a resiliency group

1 Prerequisites

- Determine the potential affect of deleting the resiliency group. What is the benefit (if any) to deleting it from Resiliency Platform management? Does this benefit outweigh the fact that the group can no longer be monitored, managed, or protected through Resiliency Platform?
- If the resiliency group is configured for disaster recovery, you cannot remove it. You must unconfigure disaster recovery before you can remove the group.
- If necessary, notify users of the upcoming change.

2 Navigate



Assets > Resiliency Groups tab

3 Select

Use the state drop-down list, **Search** field, and table heading sort feature to locate the resiliency group.

4 Remove

To remove the resiliency group, do one of the following:

- Right click on the resiliency group row and select **Delete**.
- On the right side of the resiliency group row, click on the vertical ellipsis and select **Delete**.

On the **Delete Resiliency Group** screen, click **Submit**. On the confirmation screen, click **Done**.

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 platform configuration and asset configuration information for that data center. Click **More** for detailed information and recent activity.

Resiliency Groups and Virtual Business Services summaries

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

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

Virtual Machines by Type and Platform

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

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.

Virtual Machines and Applications by Recovery Readiness

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

Use the drop-down list to filter your results.

See [“Displaying resiliency group information and status”](#) on page 23.

Understanding asset types

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

Asset	Description
Resiliency Group	A group of applications or virtual machines under Resiliency Platform control. You can use Resiliency Platform to start and stop the resiliency group, as well as protect and manage it.
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:

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

For managed 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 by Datacenter	Provides details about the resiliency groups in the data centers across all sites.
VM Inventory	Provides the platform distribution and the OS distribution details of the virtual machines that are deployed in the data centers in the form of a pie chart.
Virtual Infrastructure Inventory	Provides information about the virtual infrastructure inventory across data centers. A pie charts show the platform and virtualization technology distribution of the virtual servers across all data centers.
Migrate and Takeover	Provides a summary of the last migrate and takeover operations that were performed on the resiliency groups. A pie chart shows the percentage of successful and failed operations.
Rehearse	Provides a summary of the latest rehearse operations that were performed on the resiliency groups. A pie chart shows the percentage of successful and failed operations. A list of resiliency groups on which the rehearse operation had failed is shown. A table displays the details of the last performed rehearse operation on the resiliency groups.

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*.

Using Resiliency Platform for disaster recovery

This chapter includes the following topics:

- [About disaster recovery using Resiliency Platform](#)
- [Understanding the role of resiliency groups in disaster recovery operations](#)

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 takeover.
- Resiliency plan (a sequential execution of predefined steps) to automate site-level recovery operations on your IT infrastructure in an 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 an 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 37.

Understanding the role of resiliency groups in disaster recovery operations

To perform disaster recovery (DR) operations on virtual machines or applications, first they must be placed in a resiliency group, which is the unit of failover in Veritas Resiliency Platform.

You can configure resiliency groups without enabling them for disaster recovery. You can perform the start or stop operations on resiliency groups that are not enabled for DR. However, you cannot perform DR operations on a resiliency group without first enabling the resiliency group for disaster recovery. You can enable disaster recovery when you create the resiliency group, or at a later point of time you can select the resiliency group and perform the **Configure DR** operation.

After you enable and configure disaster recovery on a resiliency group, you can proceed with DR-specific tasks on the resiliency group, such as migrate and takeover.

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

See [“About resiliency groups”](#) on page 17.

VMware virtual machines disaster recovery

This chapter includes the following topics:

- [Understanding virtual machine disaster recovery](#)
- [Limitations for virtual machine disaster recovery](#)
- [VMware virtual machines disaster recovery - an overview of key steps](#)

Understanding virtual machine disaster recovery

A resiliency group represents a logical collection of virtual machines or applications representing a business service. Starting or stopping a resiliency group starts or stops the virtual machines or applications that are part of it. To perform disaster recovery (DR) operations, you must first enable disaster recovery for each resiliency group by performing the Configure DR operation. Once the Configure DR operation is successful, you can perform operations such as migrate, takeover, and rehearse.

See [“VMware virtual machines disaster recovery - an overview of key steps”](#) on page 40.

Limitations for virtual machine disaster recovery

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

Table 6-1

Limitations	Descriptions
Replication limitations	<p>The following are the replication-based limitations of virtual machines:</p> <ul style="list-style-type: none"> ■ Only synchronous and asynchronous mode of replications are supported in EMC SRDF. ■ Only Consistency Group (SRDF device group) based replications are supported. ■ Does not support composite group based replications and individual disk or device file based replication. ■ Raw device mapping (RDM) mapped replicated LUNs are not supported. ■ In case of Hitachi TrueCopy/Universal Replicator replication technology, the HORCM daemons must be configured and running at all times on some hosts in both the Data Centers. ■ If Hitachi Shadow Image technology is used for taking snapshots of the replicated devices, the ShadowImage pairs must be created using the -m noread option. This disables read access to the snapshot devices and helps in importing the volumes on replicated devices on to the host.
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>

VMware virtual machines disaster recovery - an overview of key steps

This section lists the key steps required to configure the disaster recovery of VMware virtual machines using Veritas Resiliency Platform.

Table 6-2 VMware virtual machines disaster recovery - an overview of key steps

Action	Description	Refer to
Set up your replication environment	Set up your VMware environment and storage arrays for replication.	<p>For EMC SRDF-based replication:</p> <p>See “Configuring VMware virtual machines disaster recovery using EMC SRDF replication” on page 77.</p> <p>For NetApp SnapMirror based replication:</p> <p>See “Configuring VMware disaster recovery using NetApp SnapMirror” on page 80.</p> <p>For EMC RecoverPoint based replication:</p> <p>See “Configuring VMware virtual machines disaster recovery using EMC RecoverPoint replication” on page 83.</p> <p>For Hitachi TrueCopy/HUR based replication:</p> <p>See “Configuring VMware virtual machines disaster recovery using Hitachi TrueCopy/ Hitachi Universal Replicator replication” on page 86.</p> <p>For HPE 3PAR Remote Copy based replication:</p> <p>See “Configuring VMware virtual machines disaster recovery using HPE 3PAR Remote Copy replication” on page 88.</p>
Add the asset infrastructure	Add the asset infrastructure to the Infrastructure Management Server (IMS) using the Resiliency Platform web console.	Refer to the <i>Veritas Resiliency Platform Deployment Guide</i> .

Table 6-2 VMware virtual machines disaster recovery - an overview of key steps *(continued)*

Action	Description	Refer to
Configure your assets for disaster recovery	Group the required virtual machines in a resiliency group and enable disaster recovery for the resiliency group.	See “Protecting virtual machines” on page 20. See “Configuring disaster recovery for a resiliency group of virtual machines” on page 45.
DR operations	Perform the required DR operations: Migrate, takeover, and rehearse.	See “Ensuring the disaster recovery readiness of your Resiliency Platform assets using the rehearse operation” on page 49. See “Performing rehearsal cleanup” on page 51. See “Migrating a resiliency group of virtual machines” on page 60. See “Taking over a resiliency group of virtual machines” on page 61.

Preparing for disaster recovery operations

This chapter includes the following topics:

- [Configuring subnet information for a data center](#)
- [Setting up subnet mapping between production and recovery data centers](#)
- [Setting up virtual switch mapping between production and recovery data centers](#)
- [Configuring disaster recovery for a resiliency group of virtual machines](#)
- [Viewing the details of a disaster recovery-enabled resiliency group](#)

Configuring subnet information for a data center

To configure subnet information for a data center

1 Navigate



Settings (menu bar)

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



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

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

You can optionally specify if a subnet should be used for the purpose of rehearsal or for production.

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

Setting up subnet mapping between production and recovery data centers

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

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

Use the **Recovery Automation** option on the Resiliency Platform web console to map your production data center's subnet with the recovery (DR) data center's subnet. Note that the subnets will be discovered only when the virtual machines are running. If a virtual machine is down at the recovery data center, subnets will not be discovered.

To set up subnet mapping between production and recovery data centers

- 1 Navigate



Recovery Automation (navigation pane)

- 2 Click **Subnet** in the **Network Mapping** page.
- 3 Click **Create Pair**.
- 4 In the **Configure Subnet Mapping - Select Subnet** page, select the subnet that should be the part of your subnet pair. You can organize the subnet into **Host** or **Datacenter** group using the **Group By** drop-down menu. Click **Next**.
- 5 In the **Configure Subnet Mapping - Pair Subnet**, select the other member of your subnet pair, and click **Submit**. The created subnet pair is listed in the **Network Mapping** page.

Setting up virtual switch mapping between production and recovery data centers

Using virtual switch mapping, you can map the virtual switch of a virtualization server at the production data center to the virtual switch of another virtualization server configured at the recovery data center.

To set up virtual switch mapping between production and recovery data centers

- 1 Navigate



Recovery Automation (navigation pane)

- 2 Click **Virtual Switch** in the **Network Mapping** page.
- 3 Click **Create Pair**.
- 4 In the **Configure Virtual Switch Mapping - Select Source Virtual Switch** page, select the virtual switch that should be the part of your virtual switch pair. Click **Next**.

- 5 In the **Configure Virtual Switch Mapping - Select destination Virtual Switch** page, select the other member of your virtual switch pair, and then click **Next**. The created virtual switch pair is listed in the **Network Mapping** page.

Note: You need to map the virtual local area network (VLAN) IDs of the primary site with the VLAN IDs of the DR site. You need to manually enter the VLAN IDs of the DR site.

- 6 In the **Configure Virtual Switch Mapping - Select VLAN pairs**, you need to manually enter the virtual local area network (VLAN) IDs of the DR data center to map the VLAN IDs of the production data center with the DR data center, and then click **Submit**.

Configuring disaster recovery for a resiliency group of virtual machines

When configuring disaster recovery (DR), Veritas 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 and applies the Subnet, VSwitch, PortGroup mappings details to the current network settings that needs to be applied after Migration in the disaster recovery data center. The Resiliency Platform stores and uses this configuration at the time of disaster recovery operations, such as, Migrate, Takeover, or Rehearse.

Note: If there are any changes to the storage stack, or network settings in any of the resiliency group members, please make sure to re-run the DR Configuration wizards so that the latest storage and network configuration snapshot are recorded.

For VMware virtual machines consider the following:

- After performing the configure DR operation, the 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.
- If a virtual machine has more than one ethernet adapters then all of them should have either static IP configuration or DHCP IP configuration. 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.

Use this procedure to configure disaster recovery (DR) for a selected resiliency group. A successful DR configuration enables takeover, migrate, and rehearse operations.

To configure disaster recovery for a resiliency group

1 Navigate



Assets (navigation pane)

Resiliency Groups

2 Double-click the desired resiliency group.

3 In the resiliency group details page, click **Configure DR**.

4 In the **Select DR Data center** page, select the target data center, and click **Next**.

See [“Configuring disaster recovery - DR data center selection”](#) on page 47.

5 The **VM Selection** page displays the matching virtual machines available at the DR data center. Click **Next**.

6 In the **Summary** page, review the information on virtual machine pairing and network customization.

See [“Configuring disaster recovery - Summary”](#) on page 47.

7 Click **Submit** to complete the disaster recovery operation for the resiliency group.

8 Post configuration, verify whether the **DR Status** column on the resiliency group details page displays the status of the resiliency group as **CONFIGURED**.

9 After you have successfully configured the resiliency group for DR operations, Resiliency Platform invokes a workflow which initializes the virtual machines for DR operations.

Ensure that this operation is successfully completed by checking in **Current** or **Completed** activities page. The operation will be listed as "**Process for initializing resiliency group for DR operations.**"

See [“Viewing activities”](#) on page 63.

Configuring disaster recovery - DR data center selection

This page lists all data centers that are currently configured in your environment. Select the appropriate data center for the disaster recovery of the resiliency group.

Table 7-1 DR data center selection panel options

Field	Description
Name	Displays the name of the disaster recovery data center.
Location	Displays the geographical location of the disaster recovery data center.

See [“Configuring disaster recovery for a resiliency group of virtual machines”](#) on page 45.

Configuring disaster recovery - Summary

This page displays the information on virtual machines pairing and the network customization for the virtual machines.

Table 7-2 Configuring disaster recovery - Summary

Field	Description
Matching Virtual Machines found in the Recovery Site	Displays the number of virtual machines at the recovery data center that match with the production data center's virtual machines. This pairing is based on virtual machine IDs.
Virtual Machines to be enabled in the Recovery Site	Displays the number of virtual machines that you need to create or register at the DR data center.
Network Customization	<p>Select this check box to apply preconfigured network settings for the virtual machines. The network customization includes the subnet and vSwitch pairing from production to recovery data center. The IP addresses for the virtual machines at the recovery data center will be applied based on the subnet mappings.</p> <p>Note: The customization is applicable only if DHCP is not configured for the data center.</p>

See [“Configuring disaster recovery for a resiliency group of virtual machines”](#) on page 45.

Viewing the details of a disaster recovery-enabled resiliency group

The Veritas Resiliency Platform console provides information about a resiliency group for which disaster recovery (DR) operation is configured successfully. The information 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.

From this view, you can also set the replication lag threshold. For more information on setting the threshold, see:

See [“Setting up replication lag threshold”](#) on page 54.

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

To view the details of a disaster recovery-enabled resiliency group

1 Navigate



Assets (navigation pane)

Resiliency Groups

2 On the resiliency groups tab, double-click the resiliency group for which disaster recovery is already configured. That is, the **DR Status** column shows the status of the resiliency group as **Configured**.

See [“Displaying resiliency group details”](#) on page 25.

Rehearsing DR operations to ensure DR readiness

This chapter includes the following topics:

- [Ensuring the disaster recovery readiness of your Resiliency Platform assets using the rehearse operation](#)
- [Rehearse operations in VMware virtual machines](#)
- [Performing rehearsal cleanup](#)

Ensuring the disaster recovery readiness of your Resiliency Platform assets using the rehearse operation

Use the **Rehearse** 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 rehearse operation

1 Navigate



Assets (navigation pane)

Resiliency Groups

2 Do one of the following on the resiliency group for which DR is already configured. That is, the **DR Status** column shows the status of the resiliency group as Configured.

- Double-click the resiliency group and select **Rehearse**.
- Click on the vertical ellipsis and select **Rehearse**.

3 Select the recovery data center and then click **Submit**.

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

See [“Performing rehearsal cleanup”](#) on page 51.

Rehearse operations in VMware virtual machines

Rehearse operations with EMC SRDF based replication:

- Device group should be associated with the snapshot LUNs. Veritas 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 operations is initiated, Veritas Resiliency Platform creates point in time snapshots as part of rehearsal operations, since it cannot work with existing snapshots.

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

- The datastores on the snapshot device are attached on the DR host.
- Veritas Resiliency Platform registers the virtual machines in the production data center for rehearsal. They have identical configuration as DR virtual machines, except these virtual machines consume storage from the datastore mounted

using the snapshot volumes. And these virtual machines will be disconnected from the network and will be unregistered during cleanup.

Rehearse operations with NetApp SnapMirror based replication:

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

Note: Rehearse operations 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.

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

See [“Ensuring the disaster recovery readiness of your Resiliency Platform assets using the rehearse operation”](#) on page 49.

Performing rehearsal cleanup

After you have performed the rehearse 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 rehearsal cleanup 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 rehearsal cleanup

1 Navigate



Assets (navigation pane)

Resiliency Groups

2 Do one of the following on the resiliency group for which DR is already configured. That is, the **DR Status** column shows the status of the resiliency group as Configured.

- Double-click the resiliency group and select **Rehearse Cleanup**.
- Click on the vertical ellipsis and select **Rehearse Cleanup**.

3 Select the data center, and then click **Submit**.

See [“Ensuring the disaster recovery readiness of your Resiliency Platform assets using the rehearse operation”](#) on page 49.

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.

See [“Setting up replication lag threshold”](#) on page 54.

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 9-1

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 9-2

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.

Setting up replication lag threshold

Veritas Resiliency Platform enables you to set up the replication lag or service level agreement (SLA) threshold.

To set up replication lag threshold

1 Navigate



Assets > Resiliency Groups tab

- 2 On the resiliency groups tab, double-click the resiliency group for which disaster recovery is already configured. The next page provides the details about the resiliency group.
- 3 Under **Replication**, enter the value for **Replication lag threshold**. Select the unit of time, and click **Save**.



See [“About risk insight”](#) on page 53.

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 9-3 Ways to display risks

To display ...	Do the following:
A complete list of risks across the resiliency domain	<ol style="list-style-type: none">1 On the menu bar, select  More Views > Risks2 On the Risk page, double-click a risk in the table to display detailed information.
Risks that are associated with a specific resiliency group or virtual business service	<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 9-4 lists the predefined risks available in Resiliency Platform. These risks are reflected in the current risk report and the historical risk report.

Table 9-4 Predefined risks

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
New VM added to replication storage	Checks if a virtual machine that is added to a consistency group on a primary site, is not a part of the resiliency group.	5 minutes	Error	<ul style="list-style-type: none"> ■ Migrate ■ Takeover ■ Rehearse 	Add the virtual machine to the resiliency group.
Replication lag exceeding threshold	Checks if the replication lag exceeds the thresholds that are defined by the user for each resiliency group.	5 minutes	Warning	<ul style="list-style-type: none"> ■ Migrate ■ Takeover 	Contact the appropriate administrator
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.
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 ■ Rehearse 	Delete or move some files or uninstall some non-critical applications to free up some disk space.
Control host not reachable	Checks if the discovery daemon is down on the Control Host.	15 minutes	Error	<ul style="list-style-type: none"> ■ Migrate 	Resolve the discovery daemon issue.

Table 9-4 Predefined risks (*continued*)

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
ESX not reachable	Checks if the ESX server is in a disconnected state.	5 minutes	Error	<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	Resolve the virtualization server connection issue. In case of a password change, resolve the password issue.
Insufficient compute resources on failover target	Checks if there are insufficient CPU resources on failover target in a virtual environment.	6 hours	Warning	<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.

[Table 9-5](#) describes some risks that are displayed in Resiliency Platform console, but these risks are not reflected in the risk reports.

Table 9-5 Other risks

Risk	Description
HOST_SFMMH_REINSTALLED	The host is disconnected. The probable cause is that the host has been reinstalled. Changes you make after this condition are not reflected on the Resiliency Manager. To correct this issue, remove and re-add this host to the Infrastructure Management Server (IMS).
HOST_DISCONNECTED_MAC_CHANGED	The host is disconnected. The probable cause is that the media access code (MAC) address of host has changed. Changes you make after this condition are not reflected on the Resiliency Manager. To correct this issue, remove and re-add this host to the Infrastructure Management Server (IMS).
VMWARE_DISCOVERY_FAILED	VMware discovery failed.
FS_FILESYSTEM_FULL	The file system is at 100% usage.

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.

Performing disaster recovery operations

This chapter includes the following topics:

- [Migrating a resiliency group of virtual machines](#)
- [Taking over a resiliency group of virtual machines](#)
- [Performing the resync operation](#)

Migrating a resiliency group of virtual machines

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

To migrate virtual machines

- 1 Navigate



Assets (navigation pane)

Resiliency Groups

- 2 Double-click the resiliency group for which DR is already configured. That is, the **DR Status** column shows the status of the resiliency group as CONFIGURED.

- 3 On the resiliency group details page, click **Migrate**.
- 4 Select the target data center and then click **Submit**.

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 any **disaster or natural calamities**, 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 updated. 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 recovered checkpoint.

To perform takeover operation on virtual machines

- 1 Navigate



Assets (navigation pane)

Resiliency Groups

- 2 Double-click the resiliency group for which DR is already configured. That is, the **DR Status** column shows the status of the resiliency group as CONFIGURED.
- 3 On the resiliency group details page, click **Takeover**.
- 4 Select the target data center, and then click **Submit**.

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. Later when the production site is up and running you need to prepare the site for next failover or migrate operation. This includes cleaning up any residue and resuming the replication from recovery to 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, deregistering virtual machines, unmounting file systems, datastores, etc.

Resync operation can be performed only if the last Takeover operation was successfully completed.

Note: Resync operation must be performed at an individual resiliency group level.

Performing the resync operation

1 Navigate



Assets (navigation pane)

Resiliency Groups

- 2** Double click the resiliency group for which DR is already configured. That is, the DR Status column shows the status of the resiliency group as **Configured**.
- 3** On the resiliency group details page, click **Resync**.
- 4** In the **Resync** panel, select the production data center name from the drop-down list, and click **Submit**.

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 63.

See [“Aborting a running activity”](#) on page 64.

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

Do one of the following:



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 64.

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 Step [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 63.

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 65.

See [“Creating a new resiliency plan template”](#) on page 66.

See [“Editing a resiliency plan template”](#) on page 70.

See [“Deleting a resiliency plan template”](#) on page 70.

See [“Viewing a resiliency plan template”](#) on page 71.

See [“Creating a new resiliency plan”](#) on page 71.

See [“Editing a resiliency plan”](#) on page 72.

See [“Deleting a resiliency plan”](#) on page 73.

See [“Executing a resiliency plan”](#) on page 73.

See [“Viewing a resiliency plan”](#) on page 74.

See [“Creating a schedule for a resiliency plan”](#) on page 74.

See [“Editing a schedule for a resiliency plan”](#) on page 75.

See [“Deleting a schedule for a resiliency plan”](#) on page 75.

See [“Viewing a schedule for a resiliency plan”](#) on page 76.

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 or virtual business service (VBS). Then you can add a resiliency group or VBS to the template to create a plan. You can create multiple plans using the same template.

You can create customized resiliency plans for the following tasks:

- Start a resiliency group.

- Stop a resiliency group.
- Migrate a resiliency group.
- Takeover a resiliency group.
- Rehearse a resiliency group.
- Clean rehearsal for a resiliency group.
- Manual task
- Run 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, rehearse, or cleanup operation, configure disaster recovery task must be successful on the selected resiliency group or VBS.

See [“Creating a new resiliency plan template”](#) on page 66.

See [“Creating a new resiliency plan”](#) on page 71.

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.
- Rehearse and rehearse cleanup of a resiliency group.
- Migrate and takeover a resiliency group.
- Manual task
See [“About manual task”](#) on page 67.
- Run a custom script
See [“About custom script”](#) on page 68.

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 65.

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 66.
See [“Creating a new resiliency plan”](#) on page 71.
- 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 1.1 or later on the Resiliency Manager, Infrastructure Management Server (IMS) and the hosts executing custom scripts. In addition, if you are using VRP together with Veritas InfoScale, the Resiliency Platform Enablement Add-on have 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 66.

See [“Creating a new resiliency plan”](#) on page 71.

- 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 66.

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 66.

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.
- Rehearse and rehearse cleanup of a resiliency group.
- Migrate and takeover a resiliency group.
- Manual task
See [“About manual task”](#) on page 67.
- Run a custom script
See [“About custom script”](#) on page 68.

Note: To create a plan for migrate, takeover, rehearse, or cleanup 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 65.

See [“Deleting a resiliency plan”](#) on page 73.

See [“Executing a resiliency plan”](#) on page 73.

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 71.

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 71.

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 71.

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 72.

See [“Deleting a resiliency plan”](#) on page 73.

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 75.

See [“Deleting a schedule for a resiliency plan”](#) on page 75.

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.

Configuring VMware environment for disaster recovery using replication

This appendix includes the following topics:

- [Configuring VMware virtual machines disaster recovery using EMC SRDF replication](#)
- [Configuring VMware disaster recovery using NetApp SnapMirror](#)
- [Configuring VMware virtual machines disaster recovery using EMC RecoverPoint replication](#)
- [Configuring VMware virtual machines disaster recovery using Hitachi TrueCopy/Hitachi Universal Replicator replication](#)
- [Configuring VMware virtual machines disaster recovery using HPE 3PAR Remote Copy replication](#)

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.

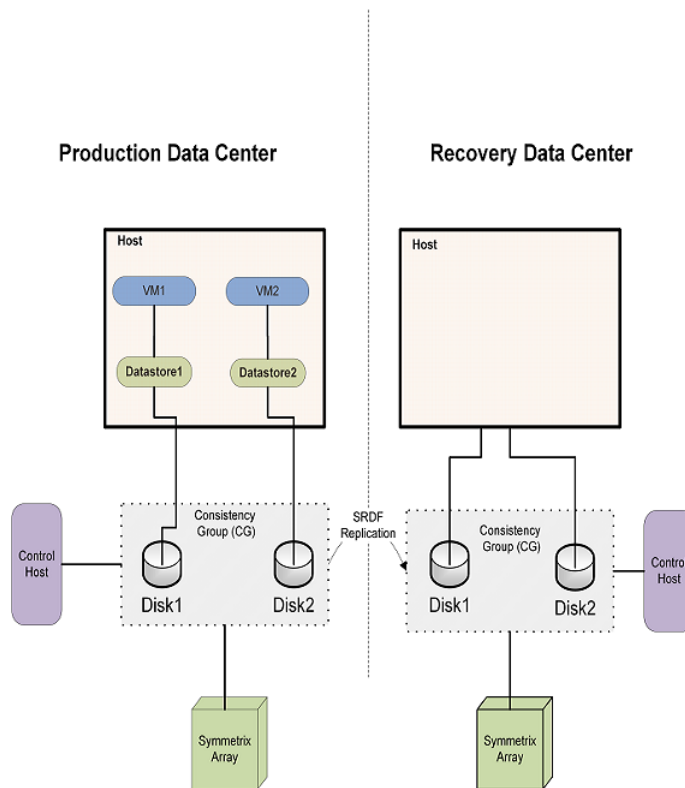
- Ensure that EMC Solutions Enabler is installed on a host and that the SRDF device groups are already set up for the replication between the primary and remote arrays.

- 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 consistency group must belong to the same resiliency group. A consistency group is a collection of Symmetrix LUNs 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).

Resiliency Platform configurations:

Using the Resiliency Platform console, you add the asset infrastructure to the IMS. The following is a summary of the steps. More information is available.

See the *Veritas Resiliency Platform Deployment Guide*.

- Add the Symmetrix enclosure to the IMS using the **Add Enclosure** option. 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 their respective IMS in each data center using the **Add Virtualization Server** option. 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 IMS using the **Add Hosts** option.
- Perform add host and add enclosure operations for the IMS at the disaster recovery data center as well.

See [“Troubleshooting discovery of assets”](#) on page 92.

Configuring VMware disaster recovery using NetApp SnapMirror

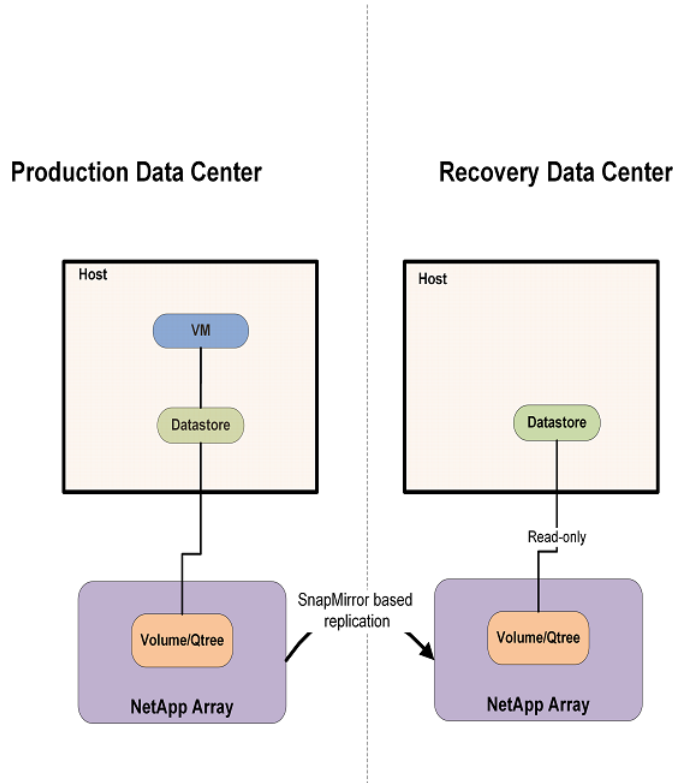
This section lists the prerequisites to enable data replication for the Veritas Resiliency Platform environment using NetApp SnapMirror when the hosts are part of a vSphere HA Cluster.

- 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. You must mount array volumes and Qtree in NFS mode only.
- 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.

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.

- Change the permission to **Grant root access to all hosts** on the replicated volumes at both the production and recovery sites.
- Ensure that all virtual machines that are part of a resiliency group have their virtual disks located on the same NAS datastores that are configured using NetApp volumes.
- Ensure that the following options are turned on for NetApp 7 Mode enclosure: `httpd.admin.enable` and `httpd.enable`.
These are required for NetApp SnapMirror operations.
- Ensure that the following licenses are installed and enabled for NetApp 7 Mode enclosure: `licensed_feature.multistore.enable` (For discovering IP addresses) and `licensed_feature.flex_clone.enable` (For rehearsal)

Note: For NetApp SnapMirror based replication in Resiliency Platform, all virtual machines that consume storage from a NetApp volume must belong to the same resiliency group. 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 VMware vCenter Servers to their respective IMS in each data center using the **Add Virtualization Server** option.. The user needs to have vCenter administrator privileges.

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

- Add the NetApp enclosure to the IMS using the **Add Enclosure** option.

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.

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

See [“Troubleshooting discovery of assets”](#) on page 92.

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.

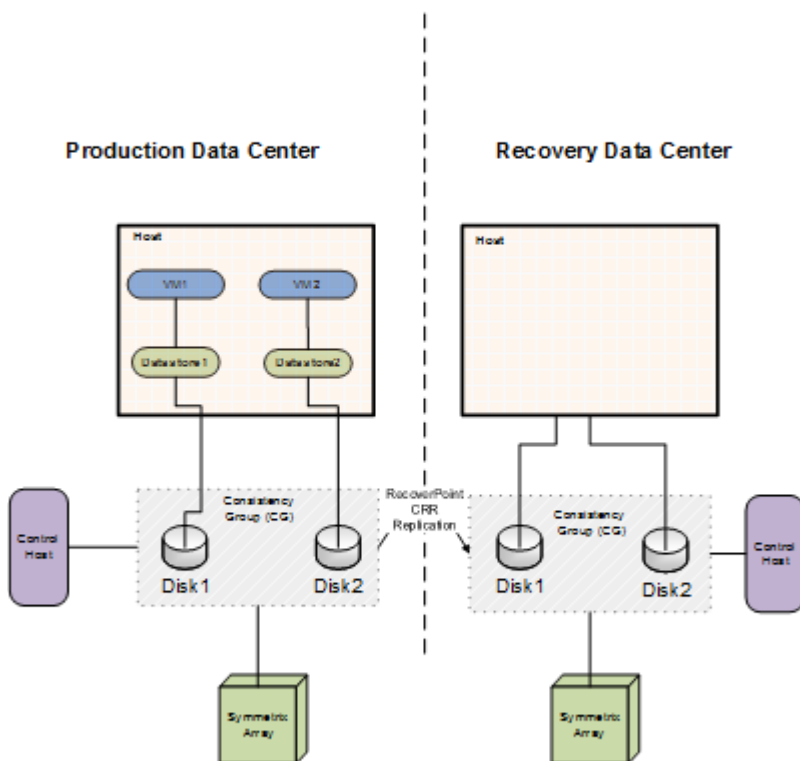
- Ensure that discovery host is able to communicate with Recoverpoint appliance using SSH.
- Confirm that RecoverPoint user has all the necessary permissions to perform EMC RecoverPoint operations.
- EMC RecoverPoint groups are set up for CRR replication between the primary and remote RecoverPoint Appliance.
- 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 consistency group must belong to the same resiliency group. A consistency group is a collection of Symmetrix LUNs 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.

Figure A-1 Configuring VMware using EMC RecoverPoint replication



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, you add the asset infrastructure to the IMS. The following is a summary of the steps. More information is available.

See the *Veritas Resiliency Platform Deployment Guide*.

- Add the backend enclosure to the IMS using the **Add Enclosure** option. 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 their respective IMS in each data center using the **Add Virtualization Server** option. The user needs to have vCenter administrator privileges.

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

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

See [“Troubleshooting discovery of assets”](#) on page 92.

Configuring VMware virtual machines disaster recovery using Hitachi TrueCopy/ Hitachi Universal Replicator replication

This section lists the prerequisites to enable data replication using Hitachi TrueCopy (HTC) or Hitachi Universal Replicator (HUR) for the Veritas Resiliency Platform environment.

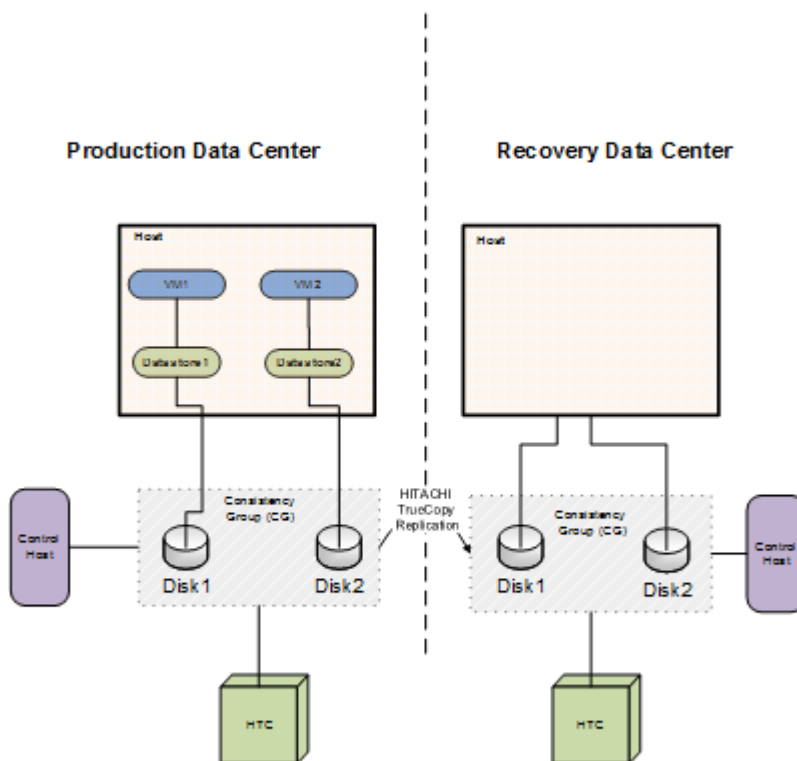
- Ensure that Hitachi Command Control Interface (HORCM CCI) is installed on a host and that the HTC Instances are already set up for the replication between the primary and remote arrays.
- If you have additionally configured Hitachi ShadowImage for taking snapshots during Rehearsals, make sure that the ShadowImage pairs are created using the -m noread flag. This ensures that the hosts cannot read the metadata from the snapshot LUNs.
- 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.

Figure A-2 Configuring VMware using Hitachi TrueCopy or Hitachi Universal Replicator replication



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 HTC enclosure to the IMS using the **Add enclosure** option. Provide the discovery host name and the HORCM location on this discovery host. Default HORCM location on Linux is `host //HORCM/usr/bin/` and default HORCM location on Windows host is `C:\\HORCM\\etc.`

Note: Any managed host can be designated as the array discovery host, including the virtual machine inside VMware ESX server that has HTC Command Device visibility and HORCM Cli installed. The host on which HTC Instance are defined can also be used as an array discovery host.

This operation returns the list of HTC 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 Add Virtualization Server option. 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 the IMS using the Add Hosts option.
- Perform add host and add enclosure operations for the IMS at the disaster recovery data center as well.

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.

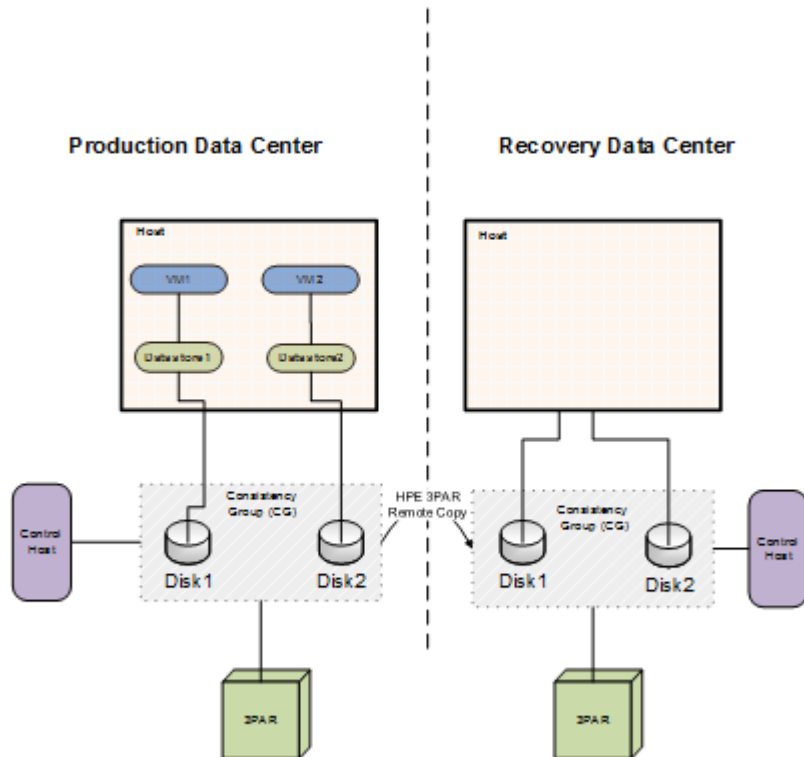
- Ensure that the discovery host is able to communicate with HPE 3PAR array using SSH.
- Confirm that HPE 3PAR array user has all the necessary permissions to perform HPE 3PAR RemoteCopy operations.

- HPE Remote Copy groups are set up for replication between the primary and remote arrays.
- 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 an 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 **Add 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 Add Virtualization Server 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.

Troubleshooting

This appendix includes the following topics:

- [Troubleshooting discovery of assets](#)
- [Viewing events and logs in the console](#)
- [Events in VMware virtual machines disaster discovery](#)

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

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 on the IMS 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 B-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 and RecoverPoint environment, the general guideline is to add/refresh the enclosure before adding/refreshing the VMware vCenter Server.

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

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

Table B-2 Configuring or refreshing asset infrastructure in IMS for VMware and EMC SRDF and RecoverPoint environment (*continued*)

Situation	Recommended sequence
You are provisioning storage from an enclosure that is already configured in the IMS to new ESX servers managed by a vCenter Server.	Refresh the enclosure first, then add the vCenter Server to the IMS or refresh it if it is already added to the IMS.

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

Table B-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*.

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.

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

- activities
 - abort 64
 - view 63
- asset infrastructure
 - troubleshooting discovery of assets 92

D

- dashboard 31
- disaster recovery
 - applications 38
 - configure 45
 - limitations 38
 - overview 40
 - resiliency group 38
 - virtual machines 38
- disaster recovery operations
 - about 36
 - migrate 60
 - rehearsal cleanup 51
 - rehearse 49
 - rehearse operations 50
 - takeover 61

E

- events 94–95

I

- Infrastructure Management Server
 - troubleshooting discovery of assets 92

L

- logs
 - viewing in console 94

M

- monitoring
 - events 95

N

- network customization 47

P

- permissions
 - about 14

R

- Recovery Automation
 - subnet mapping 43
 - virtual switch mapping 44
- replication
 - monitoring events 95
- replication lag 25
- Replication lag threshold 54
- reports
 - viewing 34
- resiliency groups
 - about 17
 - configure disaster recovery 45
 - creating from virtual machines 18, 20
 - deleting 30
 - displaying detailed information 25
 - displaying information and status 23
 - guidelines for creating 18
 - modifying 27
 - roles 37
 - starting 28
 - stopping 29
 - viewing details 48
- resiliency plan template
 - viewing 71
- resiliency plan templates
 - create 66
 - deleting 70
 - editing 70
- resiliency plans
 - about 65
 - create schedule 74
 - creating 71

resiliency plans *(continued)*

- custom script 68
- delete schedule 75
- deleting 73
- edit schedule 75
- editing 72
- executing 73
- manual task 67
- view schedule 76
- viewing 74

Resiliency Platform

- capabilities 13
- features and components 12

resync

- performing 61

risk information

- view 55

risk insight

- about 53

risks

- description 55

S

SLA threshold 54

subnet information for data center 43

V

virtual machines

- managing and monitoring using resiliency groups 18
- protecting with resiliency groups 20

VMware virtual machines disaster recovery

- using EMC RecoverPoint 83
- using EMC SRDF 77
- using NetApp SnapMirror 80