Veritas™ Resiliency Platform 2.2 Solutions for VMware

Veritas Resiliency Platform: Solutions for VMware

Last updated: 2017-07-17

Document version: Document version: 2.2 Rev 4

Legal Notice

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

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

Veritas Technologies LLC
500 E Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

Worldwide (except Japan)	CustomerCare@veritas.com
Japan	CustomerCare_Japan@veritas.com

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Section 1 Overview of Resiliency Platform 10

Chapter 1 Overview of Resiliency Platform 11

About Veritas Resiliency Platform 11
About disaster recovery using Resiliency Platform 12
About Resiliency Platform features and components 13
About Resiliency Platform capabilities 15
About managing VMware virtual machines using Resiliency Platform 16
About permissions for operations in the console 17

Chapter 2 Overview of Resiliency Platform Data Mover 18

About Veritas Resiliency Platform Data Mover 18
Supported environments for Resiliency Platform Data Mover with VMware VAIO 19
How Resiliency Platform Data Mover works 19
About Veritas Resiliency Platform Data Mover architecture 20
About full synchronization with Veritas Resiliency Platform Data Mover 22
How Veritas Resiliency Platform Data Mover handles virtual machine writes 23
Using Resiliency Platform Data Mover for recovery to premises- an overview 24

Chapter 3 Overview of recovery to on-premises data center 26

About recovery to premises using third-party replication 26
Using third-party replication for recovery to premises- an overview 27

Chapter 4 Managing assets protected by NetBackup ... 29

About NetBackup and NetBackup Appliances ... 29
About protecting assets with NetBackup using Resiliency Platform ... 30
Using NetBackup - an overview ... 31

Chapter 5 Overview of Amazon Web Services ... 32

About recovery to AWS using Resiliency Platform Data Mover ... 32
Using Resiliency Platform Data Mover for recovery to AWS- an overview ... 33

Chapter 6 Overview of vCloud ... 36

About recovery to vCloud using Resiliency Platform Data Mover ... 36
Using Resiliency Platform Data Mover for recovery to vCloud- an overview ... 37

Section 2 Preparing your environment ... 40

Chapter 7 Using array-based replication ... 41

Supported replication technologies with Veritas Resiliency Platform ... 42
Protecting VMware virtual machines using array-based replication - an overview ... 42
Configuring VMware virtual machines for disaster recovery using EMC SRDF replication ... 44
Configuring VMware virtual machines for disaster recovery using EMC RecoverPoint replication ... 46
Configuring VMware virtual machines for disaster recovery using NetApp SnapMirror ... 47
Configuring VMware virtual machines for disaster recovery using Hitachi True Copy replication ... 50
Configuring VMware virtual machines for disaster recovery using HPE 3PAR Remote Copy replication ... 51
Configuring VMware virtual machines for disaster recovery using IBM SVC Global Mirror replication ... 53
Configuring VMware virtual machines for disaster recovery using IBM XIV Remote Mirror replication ... 53

Chapter 8 Using Veritas Resiliency Platform Data Mover 55

VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover 55

Protecting VMware virtual machines using Resiliency Platform Data Mover - an overview 57

Chapter 9 Managing disaster recovery network mapping 60

Viewing and configuring network settings for a data center 60

Editing network settings for a data center 61

Removing network settings for a data center 61

Configuring DNS server settings for a data center 62

Sample command for Windows keytab file 63

Setting up network mapping between production and recovery data centers 64

Chapter 10 Managing Replication Gateway pairs 66

About Veritas Replication Gateway pairs 66

How Resiliency Platform Data Mover supports encryption for data replication 67

Creating a Veritas Replication Gateway pair 67

Modifying encryption for a Veritas Replication Gateway pair 68

Viewing Veritas Replication Gateways 68

Viewing Veritas Replication Gateway pairs 69

Removing a Veritas Replication Gateway pair 69

Section 3 Working with resiliency groups 71

Chapter 11 Managing resiliency groups 72

About resiliency groups 72

Guidelines for organizing resiliency groups 73

About service objectives 73

Managing virtual machines for basic monitoring 75

Starting a resiliency group 76

Stopping a resiliency group 77

Displaying resiliency group information and status 77

Viewing resiliency group details 80

Editing a resiliency group 80

Deleting a resiliency group 81

Chapter 12 Configuring resiliency groups for remote recovery 83

Understanding the role of resiliency groups in disaster recovery operations 84
How Resiliency Platform configures disaster recovery protection for virtual machines 84
Prerequisites for configuring VMware virtual machines for disaster recovery 85
Limitations for virtual machine disaster recovery 89
Managing virtual machines for remote recovery (DR) using 3rd party replication technology 89
Target asset selection options 91
Network customization options 92
Managing virtual machines for remote recovery (DR) using Resiliency Platform Data Mover 93
DR configuration options using Resiliency Platform Data Mover 94
Managing virtual machines for remote recovery (DR) in Amazon Web Services 96
AWS Customization options panel 97
Managing virtual machines for remote recovery (DR) in vCloud 98
Managing VMware virtual machines for remote recovery using NetBackup images 99
Verifying the replication status for Veritas Resiliency Platform Data Mover 101

Section 4 Managing disaster recovery 104

Chapter 13 Rehearsing DR operations to ensure DR readiness 105

About ensuring the disaster recovery readiness of your assets 105
Rehearse operations - array-based replication 106
Rehearse operations - Resiliency Platform Data Mover 107
Prerequisites for rehearsal operation 108
Performing the rehearsal operation 108
Performing the rehearsal operation using NetBackup images 109
Performing cleanup rehearsal 110

Chapter 14 Performing disaster recovery operations 111

How Resiliency Platform Data Mover handles DR operations 111
Migrating a resiliency group of virtual machines 113
Taking over a resiliency group of virtual machines 114
Performing the resync operation 115
Restoring data using NetBackup 116
Clearing outage 118

Chapter 15 Managing resiliency plans 120

About resiliency plans 120
Creating a new resiliency plan template 121
About manual task 122
About custom script 123
Editing a resiliency plan template 125
Deleting a resiliency plan template 125
Viewing a resiliency plan template 126
Creating a new resiliency plan 126
Editing a resiliency plan 128
Deleting a resiliency plan 128
Executing a resiliency plan 128
Viewing a resiliency plan 129
Creating a schedule for a resiliency plan 130
Editing a schedule for a resiliency plan 130
Deleting a schedule for a resiliency plan 130
Viewing a schedule for a resiliency plan 131

Chapter 16 Monitoring risks, reports, and activities 132

About the Resiliency Platform Dashboard 132
Understanding asset types 134
Displaying an overview of your assets 134
About risk insight 135
Displaying risk information 136
Predefined risks in Resiliency Platform 137
Viewing the current risk report 143
Viewing the historical risk report 144
Viewing reports 145
Managing activities 146
Viewing activities 146
Aborting a running activity 147

Chapter 17 Managing evacuation plans 149

About evacuation plan 149
Generating an evacuation plan 151
Regenerating an evacuation plan 152
Performing evacuation 153
Performing rehearse evacuation 153
Performing cleanup evacuation rehearsal 153

Appendix A General troubleshooting 155

Viewing events and logs in the console 155
Events in VMware virtual machines disaster discovery 156
Troubleshooting discovery of assets 157
Log files to troubleshoot Veritas Resiliency Platform Data Mover 159
Managing tunable parameters 160
Resiliency Platform fails to attach storage policy to virtual machines 162
Resiliency Platform fails to create storage policy 162
Resolving the Admin Wait state 162
Admin Wait state codes 163
Troubleshooting NetBackup issues 167
Troubleshooting delete resiliency group operation 168

Appendix B Sample policy and trust relationships for AWS 172

Sample policy statement for AWS 172
Sample trust relationship for AWS 173

Glossary 174

Index 176

Section 1

Overview of Resiliency Platform

- Chapter 1. Overview of Resiliency Platform
- Chapter 2. Overview of Resiliency Platform Data Mover
- Chapter 3. Overview of recovery to on-premises data center
- Chapter 4. Managing assets protected by NetBackup
- Chapter 5. Overview of Amazon Web Services
- Chapter 6. Overview of vCloud

Chapter 1

Overview of Resiliency Platform

This chapter includes the following topics:

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

About Veritas Resiliency Platform

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

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

Resiliency Platform Data Mover is a separately licensed feature of Veritas Resiliency Platform. It provides data replication between geographically separated data centers facilitating an effective disaster recovery solution. The Resiliency Platform Data Mover can be used for the following purposes:

- For recovery of VMware virtual machines to premises data center
- For recovery of VMware and Hyper-V virtual machines to cloud data center

Resiliency Platform has the following core capabilities:

Security and Compliance

Veritas Resiliency Platform provides enhanced data encryption (for data-in-flight and data-at-rest) as well as choice of data residency.

Predictability

Customers can predictably meet critical business Recovery Time Objectives (RTOs) and Recovery Point Objectives (RPOs).

Compliance

Customers can prove compliance to internal and external business continuity mandates with audit reporting and non-disruptive, real-time disaster recovery testing.

Automation

Customers get complete automation for all resiliency operations including recovery run books, and start and stop recovery orchestration for multi-tier applications. This reduces risk of downtime from human error.

Flexibility

Customers get the flexibility to keep their existing infrastructures and can innovate on their terms, with the flexibility that Resiliency Platform provides, to enable workload migration across sites and even to the cloud.

See “About Resiliency Platform features and components” on page 13.

About disaster recovery using Resiliency Platform

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

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

- Monitoring of data center assets - storage, virtual machines, and applications.
- Ability to group your virtual machines or applications in resiliency groups based on your production environment and business needs.

- Making business services more resilient by providing the ability to perform disaster recovery operations on virtual machines and applications. For example, migrate and take over.
- Ability to replicate data from virtual machines on source data centers to target data centers using Resiliency Platform Data Mover integrated with VMware API I/O filtering framework or array-based replication technologies provided by array vendors.
- Resiliency plan (a sequential execution of predefined steps) to automate site-level recovery operations on your IT infrastructure in the event of downtime.
- Auto-discovery and real-time tracking for recovery objectives.
- Ability to perform non-disruptive testing (rehearsal) on your virtual machines and applications to ensure that your infrastructure is adequately prepared for protection in the event of disaster.
- Reporting capabilities providing details about resiliency health of applications and virtual machines.

See “Understanding the role of resiliency groups in disaster recovery operations” on page 84.

About Resiliency Platform features and components

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

Resiliency Manager	The component that provides resiliency capabilities within a resiliency domain. It is composed of loosely coupled services, a distributed data repository, and a management console. The Resiliency Manager is deployed as a virtual appliance.
Infrastructure Management Server (IMS)	The component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager. The IMS is deployed as a virtual appliance. To achieve scale, multiple IMSs can be deployed in the same data center.

Veritas InfoScale Operations Manager Management Server

The component that allows discovery of InfoScale applications that are already configured in Veritas InfoScale Operations Manager. Also referred to as Veritas InfoScale Operations Manager server.

You can manage the InfoScale applications that are already configured in Veritas InfoScale Operations Manager on Linux, Solaris, AIX as well as Windows platform.

Replication Gateway

The component of Veritas Resiliency Platform Data Mover that is deployed as a virtual appliance on both data centers and used to perform replication between the data centers.

resiliency domain

The logical scope of a Resiliency Platform deployment.

It can extend across multiple data centers.

data center

For a disaster recovery use case, the resiliency domain must contain at least two data centers in different locations, a production data center and recovery data center. Each data center has a Resiliency Manager and one or more IMSs.

asset infrastructure

The data center assets that you add to Resiliency Platform for discovery and monitoring by the IMS.

The asset infrastructure can include hosts (Windows or Linux servers), virtualization servers for Hyper-V and VMware, and enclosures (storage arrays). Once the asset infrastructure is discovered by the IMS, the discovered virtual machines or applications are listed in the console as assets to manage or protect.

resiliency group

The unit of management and control in Resiliency Platform. You organize related assets into a resiliency group and manage and monitor them as a single entity.

service objective

A template to define the type of operations and technologies that are supported for a group of assets. You apply a service objective to each resiliency group.

A template which identifies the characteristics of a service. These could be availability related characteristics such as local redundancy, and number of nodes in a cluster or DR characteristics such as remote recovery, Recovery Point Objective (RPO) SLAs, rehearsal support etc. Service objective is applied when a group of assets are being added to a resiliency group.

Resiliency Platform monitors the resiliency groups based on the service objective definition and raises the risks as applicable.

Virtual Business Service (VBS)

A multi-tier business service where each VBS tier hosts one or more resiliency groups. A VBS lets you group multiple services as a single unit for visualization, automation, and controlled start and stop in the desired order. VBS uses the vertical grouping mechanism to group the multiple services.You can also perform operations such as migrate, takeover, resync, rehearsal on the entire VBS.

For more information on the above components, refer to the Deployment Guide.

About Resiliency Platform capabilities

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

Table 1-1 Resiliency Platform capabilities

Capability	More information
Configuring virtual machines and applications for remote recovery operations or basic monitoring	See “Managing virtual machines for basic monitoring” on page 75.
Starting and stopping resiliency groups for maintenance	See “Starting a resiliency group” on page 76. See “Stopping a resiliency group” on page 77.
Rehearsing disaster recovery	See “Performing the rehearsal operation” on page 108. See “Performing cleanup rehearsal ” on page 110.

Table 1-1 Resiliency Platform capabilities *(continued)*

Capability	More information
Migrating a resiliency group	See "Migrating a resiliency group of virtual machines" on page 113.
Taking over resiliency groups	See "Taking over a resiliency group of virtual machines" on page 114.
Performing the resync operation	See "Performing the resync operation" on page 115.
Restoring virtual machines from NetBackup generated backup images	See "Restoring data using NetBackup" on page 116.
Managing activities and resiliency plans	See "Managing activities" on page 146.
Displaying an overview of your resiliency domain including the number and health of your resiliency groups	See "About the Resiliency Platform Dashboard" on page 132. See "Displaying resiliency group information and status" on page 77.
Monitoring risks for protected assets	See "About risk insight" on page 135.
Viewing reports	See "Viewing reports" on page 145.

About managing VMware virtual machines using Resiliency Platform

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

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

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

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

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

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

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

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

About permissions for operations in the console

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

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

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

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

Chapter 2

Overview of Resiliency Platform Data Mover

This chapter includes the following topics:

- About Veritas Resiliency Platform Data Mover
- Supported environments for Resiliency Platform Data Mover with VMware VAIO
- How Resiliency Platform Data Mover works
- About Veritas Resiliency Platform Data Mover architecture
- About full synchronization with Veritas Resiliency Platform Data Mover
- How Veritas Resiliency Platform Data Mover handles virtual machine writes
- Using Resiliency Platform Data Mover for recovery to premises- an overview

About Veritas Resiliency Platform Data Mover

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

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

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

Features of Veritas Resiliency Platform Data Mover include the following:

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

Supported environments for Resiliency Platform Data Mover with VMware VAIO

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

- The Resiliency Platform Data Mover virtual appliances must be part of the same ESX clusters that host the protected virtual machines in the source and target data centers.
- Each host system in source and target data center must run ESXi version 6.0 U2 or later versions.
- Check the acceptance level on every host system in source and target data center. Ensure that the host system's acceptance level is not set to VMwareCertified. The allowed acceptance levels are VMwareAccepted, PartnerSupported, or CommunitySupported. By default, the ESX host is set to PartnerSupported.
- VMware tools are installed on virtual machines.
- Virtual machine disks must be in the same datastore.
- Virtual machines must not have any physical RDM disks.

How Resiliency Platform Data Mover works

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

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

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

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

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

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

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

About Veritas Resiliency Platform Data Mover architecture

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

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

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

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

- I/O receiver

  Continuously receives the virtual machine I/Os that were tapped and sent by the vtstap module in the ESXi host where virtual machines are running in a continuous fashion.
- Transceiver

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

  Applies the data to the storage after it is received on the target gateway.
- Scheduler

  Manages and schedules data transfer between gateways.
- Engine

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

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

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

Figure 2-1 Replication architecture

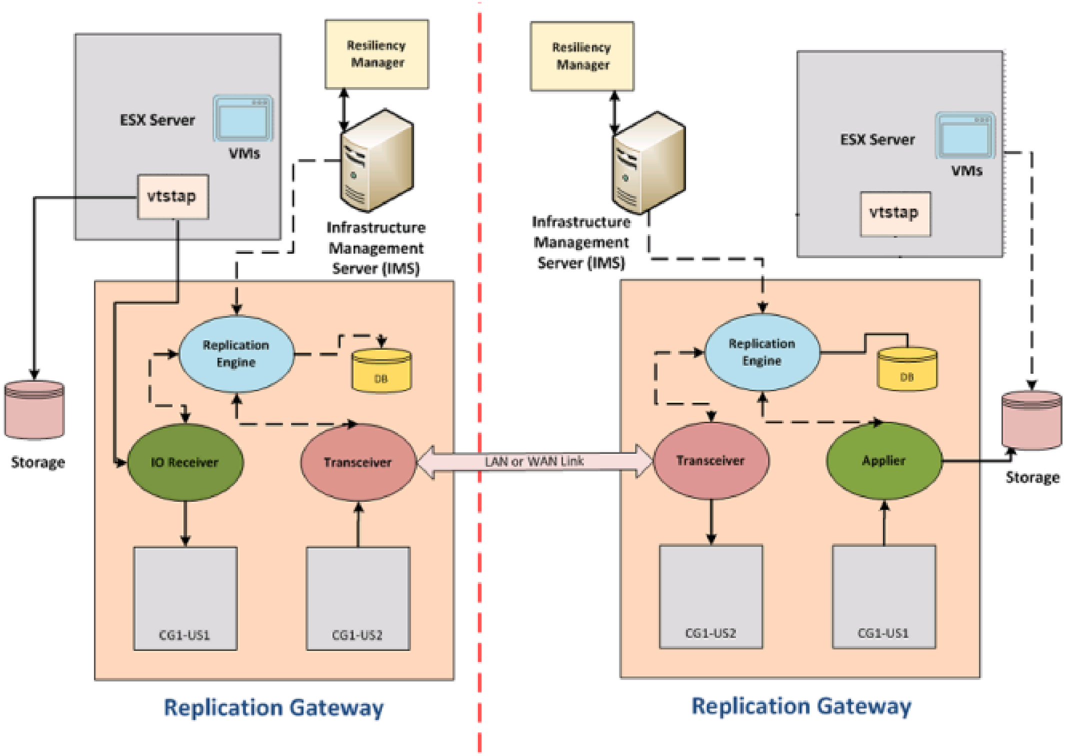

About full synchronization with Veritas Resiliency Platform Data Mover

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

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

A full synchronization is also required after a takeover operation to prepare for the next takeover or migration. Takeover is an activity initiated by a user when the source data center is down due to a disaster, and the virtual machines need to be brought up at the target (recovery) data center to provide business continuity. After a takeover, the virtual machine runs in the target (recovery) data center. Once the source (production) data center is back up and running, you must perform a Resync operation from the recovery data center before you can migrate back to the production data center. The Resync operation synchronizes the data on the production data center with the data in the recovery data center storage. When the synchronization completes, the production data center is up-to-date. You can then perform the Migrate operation.

How Veritas Resiliency Platform Data Mover handles virtual machine writes

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

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

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

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

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

Using Resiliency Platform Data Mover for recovery to premises- an overview

The following is a summary of the steps that are required to configure and protect your assets for recovery to a premises data center using the Resiliency Platform Data Mover and where to go for more information on each step.

Table 2-1 Process overview

Step	More information
Download and deploy the appropriate Resiliency Platform virtual appliances for the following components: ■ Production data center: Resiliency Manager, IMS, and Replication Gateway ■ Recovery Data center: Resiliency Manager, IMS and Replication Gateway	For more information refer to the Deployment guide.
Configure the virtual appliances as Resiliency Platform components	For more information refer to the Deployment guide.

Table 2-1 Process overview *(continued)*

Step	More information
Set up the resiliency domain using the Getting Started wizard in the web console	For more information refer to the Deployment guide.
Configure the settings for the resiliency domain	For more information refer to the Deployment guide.
Add the vCenter Server and the discovered ESX server assets to Resiliency Platform at the data centers	For more information refer to the Deployment guide.
Create gateway pair Network mapping	See "Creating a Veritas Replication Gateway pair" on page 67. See "Setting up network mapping between production and recovery data centers" on page 64.
Create resiliency groups for the virtual machines to be managed	See "Managing virtual machines for basic monitoring" on page 75. See "Managing virtual machines for remote recovery (DR) in vCloud" on page 98.
(Optional) Implement custom resiliency plans	See "Creating a new resiliency plan" on page 126.
(Optional) Configure virtual business services	For more information refer to the Solutions Guide for Virtual Business Services.
Perform disaster recovery operations.	See "Performing the rehearsal operation" on page 108. See "Performing cleanup rehearsal " on page 110. See "Migrating a resiliency group of virtual machines" on page 113. See "Taking over a resiliency group of virtual machines" on page 114. See "Performing the resync operation" on page 115.

Chapter 3

Overview of recovery to on-premises data center

This chapter includes the following topics:

- About recovery to premises using third-party replication
- Using third-party replication for recovery to premises- an overview

About recovery to premises using third-party replication

Veritas resiliency Platform provides the support for recovery of your assets to premises data center using the supported third-party replication technologies.

Following is the list of third-party replication technologies supported in Resiliency Platform:

- Array based replication:
  - EMC SRDF
  - NetApp SnapMirror
  - Hitachi True Copy/Hitachi Universal Replicator
  - EMC RecoverPoint
  - HPE 3PAR Remote Copy
  - IBM SVC Global Mirror
  - IBM XIV Remote Mirror

Using third-party replication for recovery to premises- an overview

Resiliency Platform lets you protect your data center assets and configure them for disaster recovery to on-premises data center using the supported array-based or hypervisor-based third-party replication technologies.

The following is a summary of the steps that are required to configure and protect your assets for recovery to on-premises data center and where to go for more information on each step.

Table 3-1 Process overview

Step	More information
Download and deploy the appropriate Resiliency Platform virtual appliances for the following components: ■ Production data center: Resiliency Manager and IMS ■ Recovery data center: Resiliency Manager and IMS	For more information refer to the Deployment guide.
Configure the virtual appliances as Resiliency Platform components	For more information refer to the Deployment guide.
Set up the resiliency domain using the Getting Started wizard in the web console	For more information refer to the Deployment guide.
Configure the settings for the resiliency domain	For more information refer to the Deployment guide.
Set up your replication environment	See "Protecting VMware virtual machines using array-based replication - an overview" on page 42.
Add the asset infrastructure: ■ Add Hypervisor (vCenter server, Hyper-V)	For more information refer to the Deployment guide.
Ensure that the prerequisites are met for the virtualization environment	See "Prerequisites for configuring VMware virtual machines for disaster recovery" on page 85.

Table 3-1 Process overview *(continued)*

Step	More information
Network mapping	See "Setting up network mapping between production and recovery data centers" on page 64.
Create resiliency groups for the virtual machines to be managed	See "Managing virtual machines for basic monitoring" on page 75. See "Managing virtual machines for remote recovery (DR) in Amazon Web Services" on page 96.
(Optional) Implement custom resiliency plans	See "Creating a new resiliency plan" on page 126.
(Optional) Configure virtual business services	For more information refer to the Solutions Guide for Virtual Business Services.
Perform disaster recovery operations.	See "Performing the rehearsal operation" on page 108. See "Performing cleanup rehearsal " on page 110. See "Migrating a resiliency group of virtual machines" on page 113. See "Taking over a resiliency group of virtual machines" on page 114. See "Performing the resync operation" on page 115.

Chapter 4

Managing assets protected by NetBackup

This chapter includes the following topics:

- About NetBackup and NetBackup Appliances
- About protecting assets with NetBackup using Resiliency Platform
- Using NetBackup - an overview

About NetBackup and NetBackup Appliances

About NetBackup

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

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

About NetBackup Appliances

NetBackup appliances provide a simplified solution for NetBackup configuration and the daily management of your backup environment. The appliances are rack-mount servers that run on the Linux operating system. NetBackup Enterprise

Server software is already installed and configured to work with the operating system, the disk storage units, and the robotic tape device.

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

About protecting assets with NetBackup using Resiliency Platform

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

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

See "Using NetBackup - an overview" on page 31.

Figure 4-1 Deployment architecture for NetBackup master server

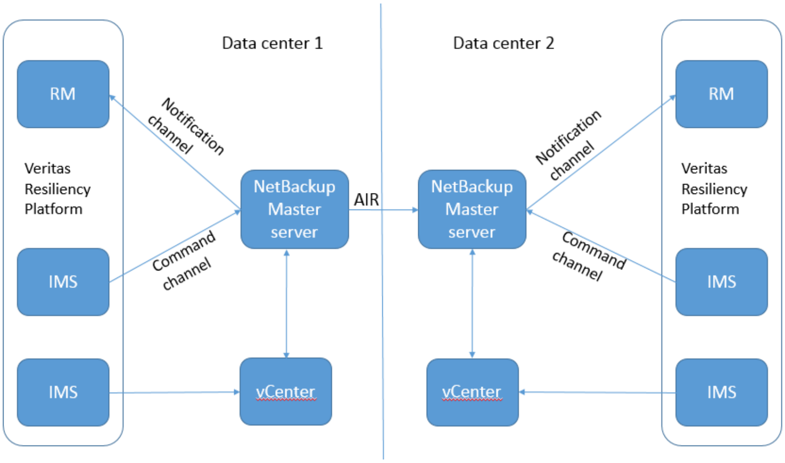

Using NetBackup - an overview

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

Table 4-1 Process overview

Step	More information
Add the NetBackup master servers to the Resiliency Manager at each data center.	For more information refer to the Deployment guide.
Add the Infrastructure Management Server (IMS) to the NetBackup master server each data center.	The IMS acts as an additional server to the NetBackup master server to discover the backup information of assets in the data center. For more information refer to the Deployment guide.
Add the VMware virtualization servers to the IMS each data center. The vCenter Server should be configured to Veritas Resiliency Platform and NetBackup with the same name or with the same IP address.	For more information refer to the Deployment guide.
Customize and activate the Copy service objective for NetBackup recovery.	For more information refer to the Deployment guide.
Customize the network mapping between the data centers.	See "Setting up network mapping between production and recovery data centers" on page 64.
Protect the assets by applying the service objective to the virtual machines.	See "Managing VMware virtual machines for remote recovery using NetBackup images" on page 99.
Restore virtual machines from backups	See "Restoring data using NetBackup" on page 116.

Chapter 5

Overview of Amazon Web Services

This chapter includes the following topics:

- About recovery to AWS using Resiliency Platform Data Mover
- Using Resiliency Platform Data Mover for recovery to AWS- an overview

About recovery to AWS using Resiliency Platform Data Mover

Veritas resiliency Platform supports recovery of your assets to AWS environment using Resiliency Platform Data Mover.

Using Veritas Resiliency Platform 2.2, you can configure and protect your VMware and Hyper-V virtual machines for recovery to AWS using the Resiliency Platform Data Mover. You will need one license for recovery and one license for Resiliency Platform Data Mover.

The following figure explains the deployment infrastructure for recovery to AWS using Resiliency Platform Data Mover:

Figure 5-1 Overview of Veritas Resiliency Platform deployment infrastructure for recovery to AWS

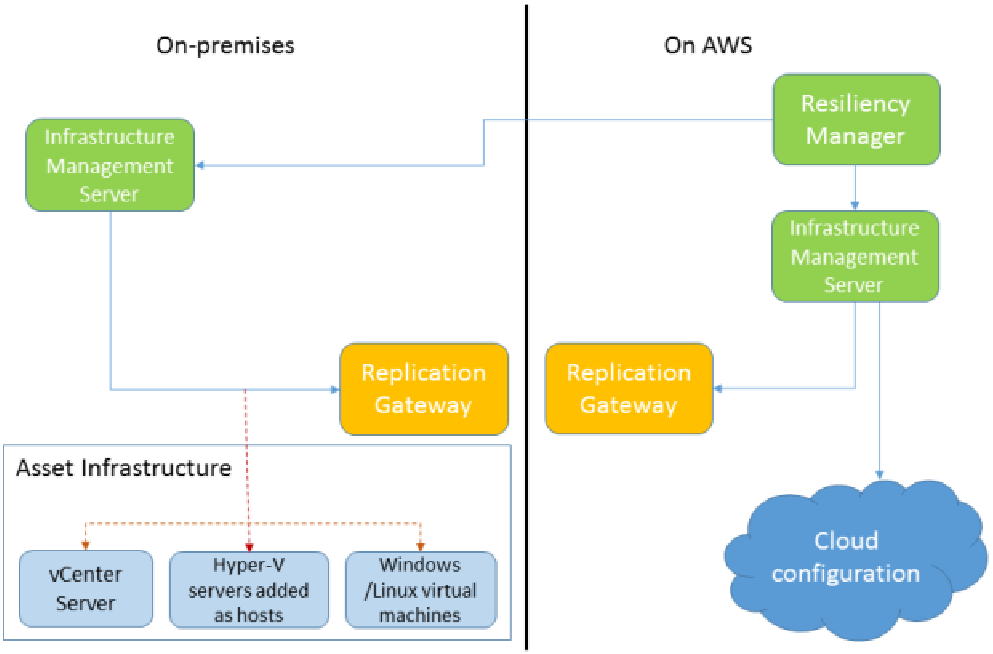

Using Resiliency Platform Data Mover for recovery to AWS- an overview

The following is a summary of the steps that are required to configure and protect your assets for recovery in Amazon Web Services (AWS) and where to go for more information on each step.

Table 5-1 Process overview

Step	More information
Download and deploy the appropriate Resiliency Platform virtual appliances for the following components: ■ In cloud: Resiliency Manager, IMS, and Replication Gateway ■ On-premises: IMS and Replication Gateway	For more information refer to the Deployment guide.
Configure the virtual appliances as Resiliency Platform components	For more information refer to the Deployment guide.
Set up the resiliency domain and add cloud configuration using the Getting Started wizard in the web console	For more information refer to the Deployment guide.
Configure the settings for the resiliency domain	For more information refer to the Deployment guide.
Add the asset infrastructure: ■ Add Hypervisor (vCenter server, Hyper-V) ■ Prepare host for replication	For more information refer to the Deployment guide.
Create gateway pair Network customization	See “Creating a Veritas Replication Gateway pair” on page 67. See “Setting up network mapping between production and recovery data centers” on page 64.
Create resiliency groups for the virtual machines to be managed	See “Managing virtual machines for basic monitoring” on page 75. See “Managing virtual machines for remote recovery (DR) in Amazon Web Services” on page 96.
(Optional) Implement custom resiliency plans	See “Creating a new resiliency plan” on page 126.
(Optional) Configure virtual business services	For more information refer to the Solutions Guide for Virtual Business Services.

Table 5-1 Process overview *(continued)*

Step	More information
Perform disaster recovery operations.	See “Performing the rehearsal operation” on page 108. See “Performing cleanup rehearsal ” on page 110. See “Migrating a resiliency group of virtual machines” on page 113. See “Taking over a resiliency group of virtual machines” on page 114. See “Performing the resync operation” on page 115.

Overview of vCloud

This chapter includes the following topics:

- About recovery to vCloud using Resiliency Platform Data Mover
- Using Resiliency Platform Data Mover for recovery to vCloud- an overview

About recovery to vCloud using Resiliency Platform Data Mover

Veritas resiliency Platform supports recovery of your assets to vCloud environment using Resiliency Platform Data Mover.

Using Veritas Resiliency Platform 2.2, you can configure and protect your VMware and Hyper-V virtual machines for recovery to vCloud using the Resiliency Platform Data Mover. You will need one license for recovery and one license for Resiliency Platform Data Mover.

The following figure explains the deployment infrastructure for recovery to vCloud using Resiliency Platform Data Mover:

Figure 6-1 Overview of Veritas Resiliency Platform deployment infrastructure for recovery to vCloud

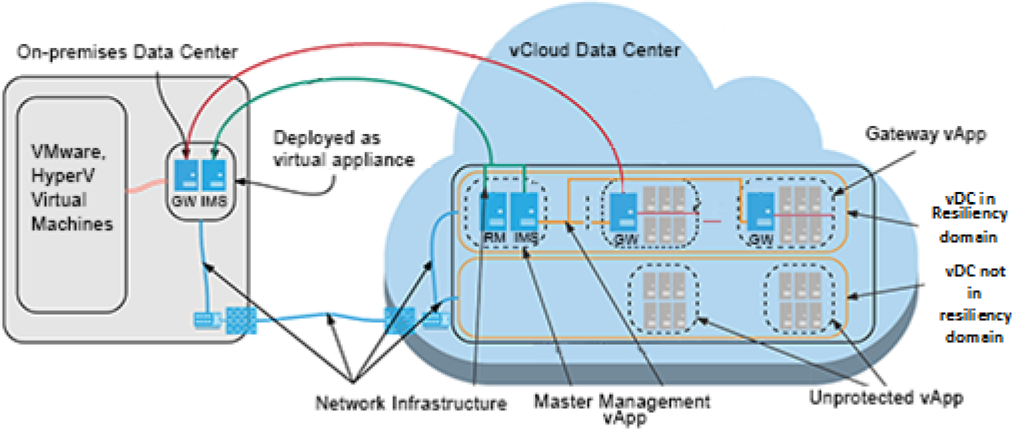

Using Resiliency Platform Data Mover for recovery to vCloud- an overview

The following is a summary of the steps that are required to configure and protect your assets for recovery in vCloud and where to go for more information on each step.

Table 6-1 Process overview

Step	More information
Steps to be performed by the cloud administrator:	
Ensure that the prerequisites are met before deploying the virtual appliances in vCloud	For more information refer to the Deployment guide.
Download the Resiliency Platform virtual appliances	For more information refer to the Deployment guide.
Upload the OVA files into catalogs	For more information refer to the Deployment guide.
Steps to be performed by the tenant:	

Table 6-1 Process overview *(continued)*

Step	More information
Deploy the appropriate Resiliency Platform virtual appliances for the following components: ■ In cloud: Resiliency Manager, IMS, and Replication Gateway If you have multiple virtual data centers, deploy Resiliency Manager and IMS in one virtual data center and only IMS in rest of the virtual data centers. ■ On-premises: IMS and Replication Gateway	For more information refer to the Deployment guide.
Configure the virtual appliances as Resiliency Platform components	For more information refer to the Deployment guide.
Set up the resiliency domain and add cloud configuration using the Getting Started wizard in the web console	For more information refer to the Deployment guide.
Configure the settings for the resiliency domain	For more information refer to the Deployment guide.
Add the asset infrastructure: ■ Add virtualization servers (vCenter server, Hyper-V server) to the on-premises data center ■ Prepare host for replication (virtual machines that you want to migrate)	For more information refer to the Deployment guide.
Create gateway pair Network mapping	See “Creating a Veritas Replication Gateway pair” on page 67. See “Setting up network mapping between production and recovery data centers” on page 64.
Create resiliency groups for the virtual machines to be managed	See “Managing virtual machines for basic monitoring” on page 75. See “Managing virtual machines for remote recovery (DR) in vCloud” on page 98.
(Optional) Implement custom resiliency plans	See “Creating a new resiliency plan” on page 126.

Table 6-1 Process overview *(continued)*

Step	More information
(Optional) Configure virtual business services	For more information refer to the Solutions Guide for Virtual Business Services.
Perform disaster recovery operations	See "Migrating a resiliency group of virtual machines" on page 113. See "Taking over a resiliency group of virtual machines" on page 114. See "Performing the resync operation" on page 115.

Section 2

Preparing your environment

- Chapter 7. Using array-based replication
- Chapter 8. Using Veritas Resiliency Platform Data Mover
- Chapter 9. Managing disaster recovery network mapping
- Chapter 10. Managing Replication Gateway pairs

Chapter 7

Using array-based replication

This chapter includes the following topics:

- Supported replication technologies with Veritas Resiliency Platform
- Protecting VMware virtual machines using array-based replication - an overview
- Configuring VMware virtual machines for disaster recovery using EMC SRDF replication
- Configuring VMware virtual machines for disaster recovery using EMC RecoverPoint replication
- Configuring VMware virtual machines for disaster recovery using NetApp SnapMirror
- Configuring VMware virtual machines for disaster recovery using Hitachi True Copy replication
- Configuring VMware virtual machines for disaster recovery using HPE 3PAR Remote Copy replication
- Configuring VMware virtual machines for disaster recovery using IBM SVC Global Mirror replication
- Configuring VMware virtual machines for disaster recovery using IBM XIV Remote Mirror replication

Supported replication technologies with Veritas Resiliency Platform

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

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

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

Protecting VMware virtual machines using array-based replication - an overview

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

Table 7-1 Configure and perform disaster recovery using array-based replication

Action	Description	Refer to
Set up your replication environment	Set up your VMware environment and storage arrays for replication	See "Configuring VMware virtual machines for disaster recovery using EMC SRDF replication" on page 44. See "Configuring VMware virtual machines for disaster recovery using NetApp SnapMirror" on page 47. See "Configuring VMware virtual machines for disaster recovery using EMC RecoverPoint replication" on page 46. See "Configuring VMware virtual machines for disaster recovery using Hitachi True Copy replication" on page 50. See "Configuring VMware virtual machines for disaster recovery using HPE 3PAR Remote Copy replication" on page 51. See "Configuring VMware virtual machines for disaster recovery using IBM SVC Global Mirror replication" on page 53. See "Configuring VMware virtual machines for disaster recovery using IBM XIV Remote Mirror replication" on page 53.
Add the asset infrastructure	Add the VMware servers and the storage arrays used for replication to Resiliency Platform	Refer to the Deployment Guide.
Prepare the virtual machines	Ensure that VMware Tools and other prerequisites are configured on virtual machines	See "Prerequisites for configuring VMware virtual machines for disaster recovery" on page 85. See "Limitations for virtual machine disaster recovery" on page 89.
Configure network settings	Configure network settings for mapping between data centers	See "Setting up network mapping between production and recovery data centers" on page 64.

Table 7-1 Configure and perform disaster recovery using array-based replication *(continued)*

Action	Description	Refer to
Configure your assets for disaster recovery	Group the virtual machines in a resiliency group and apply the appropriate service objective	See "Managing virtual machines for remote recovery (DR) using 3rd party replication technology" on page 89.
Rehearse DR operations	Test your disaster recovery environment to ensure readiness	See "Performing the rehearsal operation" on page 108. See "Performing cleanup rehearsal " on page 110.
Disaster recovery operations	Perform the disaster recovery operations	See "Migrating a resiliency group of virtual machines" on page 113. See "Taking over a resiliency group of virtual machines" on page 114. See "Performing the resync operation" on page 115.

Configuring VMware virtual machines for 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 Symmetrix LUNs composing the datastore of the virtual machine's disk are grouped together into a SRDF device group.

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

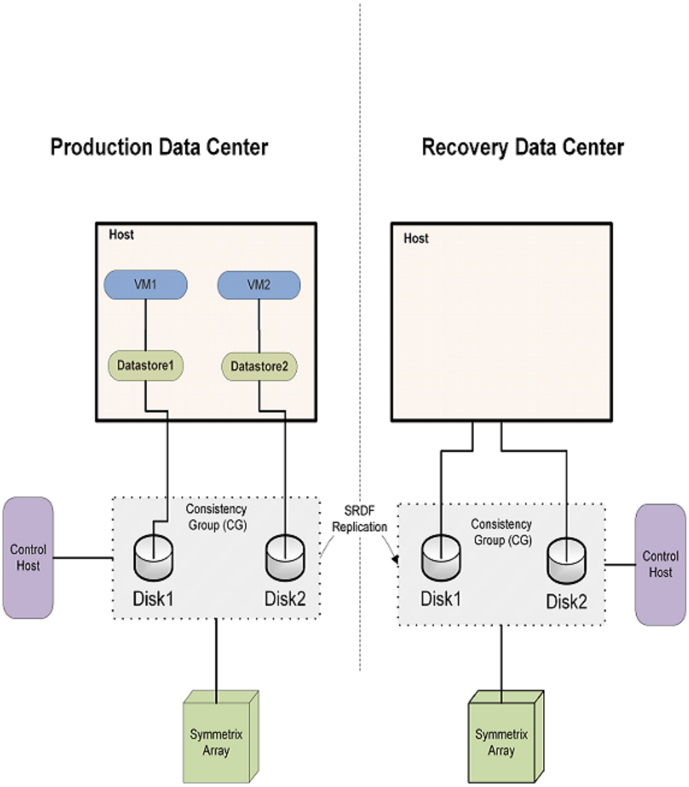

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks to add the asset infrastructure to Resiliency Platform for discovery by the Infrastructure Management Server (IMS).

Resiliency Platform configurations:

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

See the *Veritas Resiliency Platform Deployment Guide*.

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

Default SymCLI location on Linux host	/opt/emc/SYMCLI/bin/
Default SymCLI location on Windows host	C:\Program Files\EMC\SYMCLI\bin

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

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

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

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

  Ensure that the virtualization server and ESX server are discovered successfully.
- Add the host where the SRDF device groups are configured to the appropriate data centers.

See “Troubleshooting discovery of assets” on page 157.

Configuring VMware virtual machines for 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 the Infrastructure Managment Server (IMS) is able to communicate with RecoverPoint appliance using SSH.
- Ensure that EMC RecoverPoint Appliance user has *admin* role to perform EMC RecoverPoint operations.
- 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 Symmetrix LUNs composing the datastore of the virtual machine's disk are grouped together into a RecoverPoint consistency group.

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

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

Resiliency Platform configurations:

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

See the *Veritas Resiliency Platform Deployment Guide*.

- Ensure that the enclosure discovery is complete before proceeding with adding the VMware vCenter Servers
- Add the vCenter Servers to the appropriate data centers. The user needs to have vCenter administrator privileges.
  Ensure that the virtualization server and ESX server are discovered successfully.
- Add EMC RecoverPoint appliance.

See “Troubleshooting discovery of assets” on page 157.

Configuring VMware virtual machines for disaster recovery using NetApp SnapMirror

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

- Ensure that the NetApp volumes are already set up for replication between the primary and remote NetApp storage systems, and the replication has a replication schedule associated with it.
  Resiliency Platform does not support one NetApp volume having more than one SnapMirror destination volumes.

- NetApp volumes or qtrees can be provisioned as NFS datastores mounted on ESX servers, or LUNs created on NetApp volumes or qtrees can be provisioned to ESX servers via Fibre Channel and mounted as VMFS datastores.
- Ensure that the NetApp SnapMirror replicated volumes are mounted on the respective VMware ESX servers in both the sites. Do not mount the replicated peer NetApp volumes to the same VMware ESX server.

  Also ensure that the volumes are replicated using SnapMirror policy type mirror. SnapMirror policy types vault and mirror-vault are not supported.
- NFS datastore need not be mounted on the recovery data center. While creating a resiliency group, you need to select the ESX server or the cluster on the recovery data center on which you want the Resiliency Platform to mount the datastore.

  If you have already mounted the datastore then ensure that they have the same name on both data centers.
- 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.

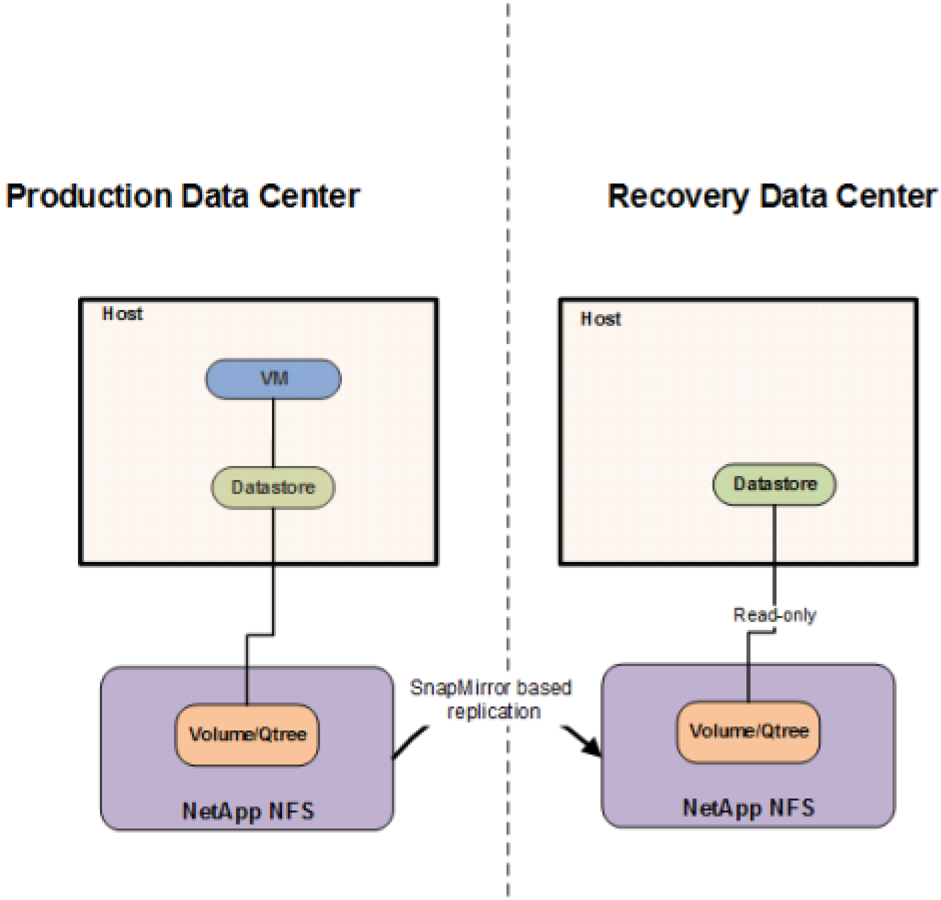

Resiliency Platform configurations:

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

For more information, see the *Veritas Resiliency Platform Deployment Guide*.

Resiliency Platform configurations:

- Configure the VMware vCenter Server to send traps to the IMS.
- Add the VMware vCenter Servers to their respective data centers. The user needs to have vCenter administrator privileges.
  Ensure that the virtualization server and ESX server are discovered successfully.
- Add the NetApp enclosures to their respective data centers.

Ensure that you have sufficient privileges to perform SnapMirror replication operations.

See “Troubleshooting discovery of assets” on page 157.

Configuring VMware virtual machines for disaster recovery using Hitachi True Copy replication

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

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

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

- Ensure that the Symmetrix LUNs composing the datastore of the virtual machine's disk are grouped together into a HTC.

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.

Resiliency Platform configurations:

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

For more information see the *Veritas Resiliency Platform Deployment Guide*.

Resiliency Platform configurations:

- Configure the VMware vCenter Server to send traps to the IMS.

- Add the HTC discovery hosts to the respective data centers. HTC discovery hosts are hosts on which HORCM package is installed and HTC device groups are created.
- Add the vCenter Servers to their respective data centers. The user needs to have vCenter administrator privileges.

  Ensure that the virtualization server and ESX server are discovered successfully.
- Add the host where the HTC Instances are configured to their respective data centers.

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

This section lists the prerequisites and limitations to enable data replication using HPE 3PAR Remote Copy replication.

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

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

- Ensure that the Symmetrix LUNs composing the datastore of the virtual machine's disk are grouped together into a 3PAR consistency group.

Note: For HPE 3PAR Remote Copy based replication, all virtual machines that consume storage from a 3PAR Remote Copy Instance must belong to the same resiliency group. An Instance is a collection of volume groups that helps in maintaining write consistency during replication.

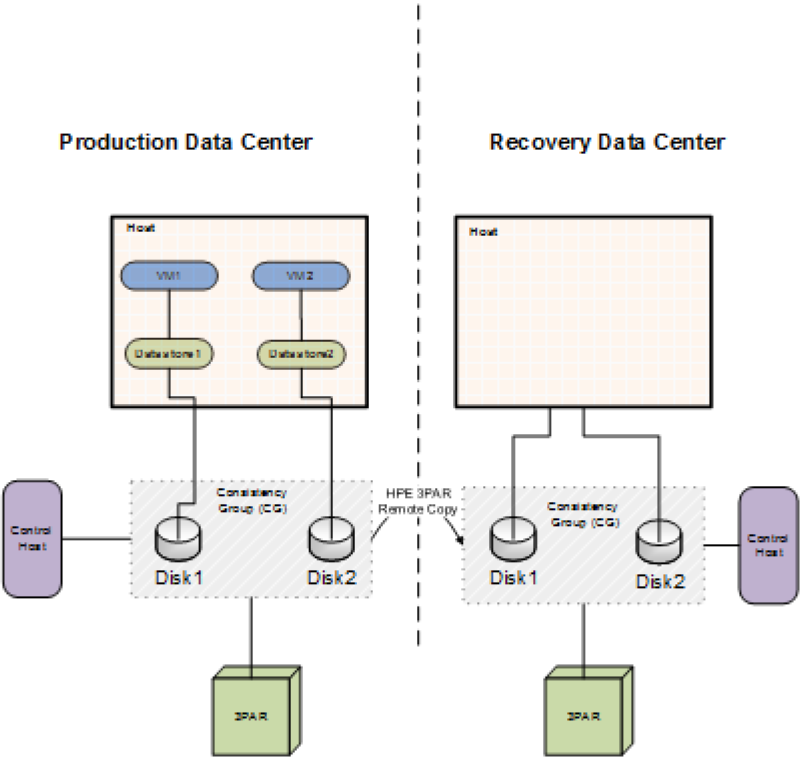

Once you have performed the necessary configurations, proceed with Resiliency Platform specific tasks to add the asset infrastructure to the Infrastructure Management Server (IMS). The following is a summary of the steps.

For more information see the *Veritas Resiliency Platform Deployment Guide*.

Resiliency Platform configurations:

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

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

- Add the vCenter Servers to their respective IMS in each data center using the **+ vCenter** option. The user needs to have vCenter administrator privileges. Ensure that the virtualization server and ESX server are discovered successfully.

Limitations:

- HPE 3PAR Remote Copy synchronous replication is not supported.
- 3PAR storage connectivity via iSCSI is not supported.

Configuring VMware virtual machines for disaster recovery using IBM SVC Global Mirror replication

Ensure that the IBM SVC Global Mirror replicated LUNs are assigned to the respective VMware ESX Servers.

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

Veritas Resiliency Platform configurations

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

For more information, refer to the *Deployment Guide*.

- Add IBM SVC enclosures to the appropriate data centers using the **+ IBM Enclosure** option. Provide the IBM SVC server, username, and password.
- Add the vCenter Servers to the appropriate data centers. The user needs to have vCenter administrator privileges. Ensure that the virtualization server and ESX server are discovered successfully.

Configuring VMware virtual machines for disaster recovery using IBM XIV Remote Mirror replication

Ensure the following to enable data replication using IBM XIV Remote Mirror.

- Ensure that the IBM XIV Consistency groups are properly configured with mirrored replication between the primary and secondary storage systems.

  Resiliency Platform does not support the Consistency groups that are not mirrored configured.

- If IBM XIV volumes are used for replication, ensure that they are properly mirror configured between the primary and secondary storage systems.
- Ensure that the IBM XIV replicated volumes are visible on the respective VMware ESX servers on both the data centers.

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.

Complete the following tasks before you proceed with the Resiliency Platform specific tasks.

- Install IBM XIV Command Line Interface (XCLI) on a Windows or Linux host. This host acts as a discovery host for array IBM XIV enclosures.
- Using the XCLI command, verify if the IBM XIV array is accessible and command execution is successful.
  For more information, refer to the *Deployment guide*.
- Ensure IMS is reachable from the host having IBM XCLI installed.
- Ensure that the IBM XIV array user is having the role of an 'Administrator'.

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

Veritas Resiliency Platform configurations

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

For more information, refer to the *Deployment guide*.

- Add the host having XCLI installed as a discovery host.
  To add a Windows host as a discovery host, you need to configure a Windows control host.
- Add IBM XIV enclosures to the appropriate data centers using the **+ IBM Enclosure** option. Select the discovery host, provide the enclosure name or IP address, username, and password. Enter the location of the XCLI binary on the discovery host.
  Ensure that the enclosures are discovered successfully.
- Add the vCenter Servers to the appropriate data centers. The user needs to have vCenter administrator privileges. Ensure that the virtualization server and ESX server are discovered successfully.

Chapter 8

Using Veritas Resiliency Platform Data Mover

This chapter includes the following topics:

- VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover
- Protecting VMware virtual machines using Resiliency Platform Data Mover - an overview

VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover

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

Table 8-1 VMware vCenter Server privileges required for Resiliency Platform Data Mover

Category	Privilege
System	System.View System.Anonymous System.Read

Table 8-1 VMware vCenter Server privileges required for Resiliency Platform Data Mover *(continued)*

Category	Privilege
Host	Host.Config.Patch Host.Cim.CimInteraction Host.Config.Maintenance Host.Config.Storage Host.Config.Settings Host.Config.Network
Virtual machine configuration	VirtualMachine.Config.AddExistingDisk VirtualMachine.Config.AddNewDisk VirtualMachine.Config.RemoveDisk VirtualMachine.Config.Rename VirtualMachine.Config.CPUCount VirtualMachine.Config.Memory VirtualMachine.Config.EditDevice VirtualMachine.Config.DiskExtend
Virtual machine interaction and inventory	VirtualMachine.Interact.PowerOn VirtualMachine.Interact.PowerOff VirtualMachine.Interact.ToolsInstall VirtualMachine.Inventory.Create VirtualMachine.Inventory.Register VirtualMachine.Inventory.Unregister VirtualMachine.Inventory.Remove
Virtual machine provisioning	VirtualMachine.Provisioning.CloneVirtualMachine VirtualMachine.Provisioning.Customize
Network privileges	Network.Assign Network.Configure
Storage policy	VM storage policies.Update VM storage policies.View Profile.create

Table 8-1 VMware vCenter Server privileges required for Resiliency Platform Data Mover *(continued)*

Category	Privilege
Folder	Folder.delete
Datastore	Datastore.AllocateSpace Datastore.FileManagement Datastore.Browse
Alarm	Alarm.Create Alarm.Delete Alarm.Modify
vCenter	Global.Settings Global.Diagnostics
Snapshots	VirtualMachine.State.CreateSnapshot VirtualMachine.State.RemoveSnapshot
Virtual appliance deployment	VApp.Import

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

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

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

Table 8-2 Configuring disaster recovery of VMware virtual machines using Data Mover

Action	Description	Refer to
Deploy the Replication Gateways	Deploy and configure the Veritas Replication Gateway virtual appliances on both data centers	Refer to the Deployment Guide.
Create the gateway pairs	Configure the Replication Gateway pairs to be used for replication	See "Creating a Veritas Replication Gateway pair" on page 67.
Add the VMware asset infrastructure	Prepare the VMware virtualization servers for the Resiliency Platform environment Add the VMware virtualization servers to Resiliency Platform at both data centers	Refer to the Deployment Guide.
Prepare the virtual machines	Ensure that VMware Tools, storage requirements for Replication Gateway and other prerequisites are configured on virtual machines	See "Prerequisites for configuring VMware virtual machines for disaster recovery" on page 85. See "Limitations for virtual machine disaster recovery" on page 89.

Table 8-2 Configuring disaster recovery of VMware virtual machines using Data Mover *(continued)*

Action	Description	Refer to
Configure your assets for disaster recovery	Group the required virtual machines in a resiliency group and choose the appropriate service objective to configure disaster recovery for the resiliency group	See “Managing virtual machines for remote recovery (DR) using Resiliency Platform Data Mover” on page 93.
Rehearse DR operations	Test your disaster recovery environment to ensure readiness	See “Performing the rehearsal operation” on page 108. See “Performing cleanup rehearsal ” on page 110.
Perform DR operations	Migrate or take over your virtual machines and resync the replicated data	See “Migrating a resiliency group of virtual machines” on page 113. See “Taking over a resiliency group of virtual machines” on page 114. See “Performing the resync operation” on page 115.

Chapter 9

Managing disaster recovery network mapping

This chapter includes the following topics:

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

Viewing and configuring network settings for a data center

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

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

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

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

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

To configure network settings for a data center

1 Navigate

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

Expand the data center > **Access Profile**

Click **+ Add Subnet**.

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

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

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

Editing network settings for a data center

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

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

To edit the network settings for a data center

1 Navigate

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

Expand the data center > **Access Profile**

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

2 Edit the name and the purpose as required.

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

Removing network settings for a data center

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

To remove the network settings for a data center

1 Navigate

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

Expand the data center > **Access Profile**

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

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

Configuring DNS server settings for a data center

Using the Resiliency Platform console, you can configure the DNS settings for the data center. You can add DNS servers for the data center or remove the settings for servers that were previously added.

Windows DNS: command to generate the keytab file and the privileges required:

- Command to generate keytab file:

```
ktpass.exe -princ <Principal Name> -mapuser <User Account>
 -pass <Password> -crypto <Encryption Type> -ptype <Principal Type>
 -out <Name for Keytab File>
```

- Ensure that you have the required privileges in Windows DNS to update Forward and Reverse Lookup zones. Refer to the Microsoft documentation for more details.

Linux Bind: command to generate private key file and the privileges required:

- Command to generate key and private file:

```
dnssec-keygen -a <Algorithm> -b <Keysize> -n HOST <Name Type>
```

- Ensure that you have the required privileges in Bind to update Forward and Reverse Lookup zones. Refer to Linux documentation for more details.

See “Sample command for Windows keytab file” on page 63.

To configure DNS server settings for a data center

1 Prerequisites

- Ensure that ports listed for DNS server are open for communication. For a list of ports to be opened on DNS server, see:
- You must have the following information:

   - The IP address of the DNS server
   - The name of the domain, and associated credentials.

     Linux Bind: For TSIG authentication, you need the TSIG key and TSIG private files.

     Windows DNS: For GSSAPI authentication, you need the user name and keytab file.
   - A test host name and IP address for performing a test operation. The test operation validates the specified DNS configuration.

2. Navigate

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

   Expand the data center > **Access Profile**

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

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

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

Sample command for Windows keytab file

Following is a sample command for Windows keytab file.

Authentication domain (AD) user is **user2**, configured on **VRPWINDNS.COM** domain. Password of the user is **user@123**. Ensure that the domain name is always in capital letters.

In the command, *princ* is the user name, *mapuser* is the user account.

Principal type (ptype), needs to be specified as *KRB5_NT_PRINCIPAL*. And *out* is the output keytab file, which is *C:/user2.keytab* for the sample.

Using the above values, the sample command is:

```
C:\Users\Administrator>ktpass.exe /princ user2@VRPWINDNS.COM /mapuser user2@VRPWINDNS.COM
 /pass user@123 /ptype KRB5_NT_PRINCIPAL /out C:/user2.keytab
```

Verifying the keytab file

After the keytab file is generated, copy the keytab file to a UNIX machine having kinit utility.

Verify the connection with DNS using - *kinit user@DOMAIN* which is *kinit user2@VRPWINDNS.COM* as per above sample values.

Enter the password of user2. On successful execution of the command, verify the keytab file using: *kinit user2@VRPWINDNS.COM -k -t /root/user2.keytab*

Setting up network mapping between production and recovery data centers

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

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

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

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

This is not applicable if the recovery data center is Amazon Web Services (AWS). For AWS, subnet to subnet mapping is sufficient.

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

If the recovery data center is AWS, then ensure that the production subnet and the rehearsal subnet are in the same virtual private cloud (VPC).

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

To set up network mapping between production and recovery data centers

1 Navigate

   Disaster Recovery Settings (navigation pane)

   Do one of the following:

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

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

2 In the **Network Mapping** page, select the source and the target data centers, and the network types that should be the part of your network pair.

   If recovery is in AWS, select subnets.

3 Click **Choose selected** or drag and drop the selections in the drag area at the bottom.

4 Click **Next** to submit your selections.

5 To remove a network pair, right-click the pair and select **Delete Pair**.

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

Chapter 10

Managing Replication Gateway pairs

This chapter includes the following topics:

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

About Veritas Replication Gateway pairs

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

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

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

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

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

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

How Resiliency Platform Data Mover supports encryption for data replication

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

Creating a Veritas Replication Gateway pair

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

To create a Replication Gateway pair

1. Prerequisites

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

2. Navigate

   Disaster Recovery Settings (navigation pane)

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

3. Specify the information in the wizard:

   - Select the gateways to be paired. You can filter each list by data center.
   - Enter the IP address to be used by the gateway. You can specify different IP addresses for communication between gateways at the source and target data centers and for communication between an ESXi host and the gateway.

- Optionally edit the default name of the gateway pair.
- Optionally change the data encryption scheme selection.

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

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

See “Removing a Veritas Replication Gateway pair” on page 69.

See “About Veritas Replication Gateway pairs” on page 66.

Modifying encryption for a Veritas Replication Gateway pair

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

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

To modify a Replication Gateway pair

1 Navigate

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

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

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

See “About Veritas Replication Gateway pairs” on page 66.

Viewing Veritas Replication Gateways

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

To view Replication Gateways

1. Navigate

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

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

2. Click **Data Mover**

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

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

See “About Veritas Replication Gateway pairs” on page 66.

Viewing Veritas Replication Gateway pairs

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

To view Replication Gateway pairs

- Navigate

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

See “About Veritas Replication Gateway pairs” on page 66.

Removing a Veritas Replication Gateway pair

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

To remove a Replication Gateway pair

1. Prerequisites

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

2. Navigate

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

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

See “About Veritas Replication Gateway pairs” on page 66.

Section 3

Working with resiliency groups

- Chapter 11. Managing resiliency groups
- Chapter 12. Configuring resiliency groups for remote recovery

Chapter 11

Managing resiliency groups

This chapter includes the following topics:

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

About resiliency groups

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

For example, you can organize several virtual machines into a resiliency group, and name it `VM_Group`. When you perform an action on `VM_Group` from the Resiliency Platform console, all the virtual machines in the group are included. For example, if you start `VM_Group`, all the virtual machines in the group start, similarly when you stop `VM_Group` all assets stop.

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

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

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

See "About service objectives" on page 73.

See "Managing virtual machines for basic monitoring" on page 75.

Guidelines for organizing resiliency groups

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

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

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

About service objectives

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

The local and remote recovery service objective includes tunables such as Recovery Point Objective (RPO) for assets being managed in that resiliency group and you would be required to select the recovery data center.

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

Following are the types of service objective templates:

- Remote recovery of applications - provides recovery operations as well as the start and stop operations for applications.
- Remote recovery of hosts - provides recovery operations as well as the start and stop operations for hosts.
- Monitor assets - provides only monitoring, that is start and stop operations.

For virtual machines you have the following two options for data availability.

- Copy: The available technology is NetBackup. This option is available only for VMware virtual machines.
  This option is available only if the acceptable RPO is 240 minutes (4 hours) and above.
- Replication: The available technologies are SnapMirror, SRDF, VRP Data Mover, RemoteCopy 3PAR, RecoverPoint, Hyper-V Replication, and Hitachi True Copy.

Note: Authorization to activate a template and edit the settings depends on the permissions that are assigned to users and groups in Resiliency Platform.

Following is the list of pre-activated service objectives:

- Recover hosts
- Recover applications
- Monitor assets
- Local and remote recovery of hosts
- Local recovery of hosts

You can view the details of both the activated service objectives and the templates in the web console. You can also delete any pre-activated service objective that you do not want to use in your environment, provided that it is not in use by any resiliency group.

The default pre-activated service objectives do not monitor an RPO. If you need RPO monitoring, activate a service objective template by providing the relevant RPO value.

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

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

Managing virtual machines for basic monitoring

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

There are two types of pre-activated service objectives:

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

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

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

To manage virtual machines for basic monitoring

1 Prerequisites

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

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

2 Navigate

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

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

3 Select the virtual machines:

- Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.
- Drag and drop virtual machines to **Selected Instances**.

For VMware assets, if you plan to configure disaster recovery protection later using Resiliency Platform Data Mover, each resiliency group must map to only one ESX cluster.

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

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

6 Supply a name for the resiliency group.

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

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

Starting a resiliency group

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

To start a resiliency group

1 Navigate

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

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

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

You can also perform operations from the Details page.

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

If you have applied update 2.0.0.100 on Veritas Resiliency Platform 2.2, you can select the checkbox on the **Start Resiliency Group** wizard to start the post-replication operations of migrate or takeover workflow on the production data center such as refreshing storage, network, compute, and customization.

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

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

Stopping a resiliency group

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

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

To stop a resiliency group

1 Prerequisites

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

2 Navigate

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

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

4 ⋮ On the row for the resiliency group, select the vertical ellipsis > **Stop**.

   You can also perform operations from the Details page.

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

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

Displaying resiliency group information and status

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

Table 11-1 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 132.
Assets > **Resiliency Groups** tab	Medium. Lists all your resiliency groups in one place.	Seeing what is in each of your data centers, the state of the groups, and so on.
Resiliency group-specific screen	Highest. Lists each asset in the resiliency group, their type, and state.	Getting detailed information on a resiliency group and its underlying assets, including disaster recovery status. This screen lists available operations for the group. See "Viewing resiliency group details" on page 80.

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

To display resiliency group information and status

1 Navigate

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

2 Review information and status

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

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

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

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

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

Viewing resiliency group details

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

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

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

To view details of a resiliency group

1 Navigate

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

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

3 ⋮ On the row for the resiliency group, select the vertical ellipsis > **Details**.

   You can also double-click the row to view details.

   The details page includes the following:

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

See “Displaying resiliency group information and status” on page 77.

Editing a resiliency group

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

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

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

A risk is also raised when you add or remove a virtual machine that belonged to DR protected resiliency group. To clear the risk, you need to edit the resiliency group and add or remove the virtual machine.

To edit the resiliency group information

1. Navigate

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

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

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

   You can also edit the resiliency group from its Details page.

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

Deleting a resiliency group

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

If the resiliency group was configured for protection using Resiliency Platform Data Mover replication, then Resiliency Platform Data Mover is unconfigured before the resiliency group is deleted. During the delete operation you can choose to delete the disks on the production data center and also choose to ignore any subtasks that fail. If you choose to ignore the failed subtasks, you need to fix them manually. Resiliency groups can be deleted from production data center, on-premises recovery data center, or from cloud recovery data center.

To successfully complete the delete operation ensure the following:

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

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

- During the operation, replication is stopped and Veritas Replication Sets are deleted on gateways and on-premises virtual machines.

- Journal disks are removed from the virtual machines on the production data center and cloud virtual machines instances are deleted.
- All the cloud virtual machines disks that are attached to the cloud Replication Gateway are deleted.

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

If you are deleting a resiliency group that was configured for protection using Resiliency Platform Data Mover replication and is active on cloud, then you can select the following options:

- During the delete operation you can choose to delete the disks on the production data center.

  If the check box is not selected, then you need to manually identify the disk, having the vxtap kernel module, which is attached to the replication gateway and delete it.
- You can choose to ignore any subtasks that fail during the delete operation.

  If you choose to ignore the failed subtasks, then you need to fix them manually.

To delete a resiliency group

1. Navigate

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

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

3. ⋮ On the row for the resiliency group, select the vertical ellipsis > **Delete**.

   You can also perform operations from the Details page

4. Confirm the deletion.

Chapter 12

Configuring resiliency groups for remote recovery

This chapter includes the following topics:

- Understanding the role of resiliency groups in disaster recovery operations
- How Resiliency Platform configures disaster recovery protection for virtual machines
- Prerequisites for configuring VMware virtual machines for disaster recovery
- Limitations for virtual machine disaster recovery
- Managing virtual machines for remote recovery (DR) using 3rd party replication technology
- Managing virtual machines for remote recovery (DR) using Resiliency Platform Data Mover
- Managing virtual machines for remote recovery (DR) in Amazon Web Services
- Managing virtual machines for remote recovery (DR) in vCloud
- Managing VMware virtual machines for remote recovery using NetBackup images
- Verifying the replication status for Veritas Resiliency Platform Data Mover

Understanding the role of resiliency groups in disaster recovery operations

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

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

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

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

How Resiliency Platform configures disaster recovery protection for virtual machines

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

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

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

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

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

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

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

Prerequisites for configuring VMware virtual machines for disaster recovery

Before you run the wizard to configure disaster recovery protection for a resiliency group of VMware virtual machines, ensure that you have met the following prerequisites for the virtualization environment:

- The vCenter and ESX servers for the virtual machines must be added to Resiliency Platform at the production and recovery data center.

  For more information on adding asset infrastructure, refer to the *Deployment Guide*.
- VMware Tools must be installed on the virtual machines.
- If you add new disks, ensure that they are visible from the guest operating system.
- If a virtual machine has more than one ethernet adapter, then all of them should have either static IP configuration or DHCP IP configuration. A mix of static and DHCP IP configuration is not supported on the same virtual machine.

- All the virtual disks must be connected to the virtual SCSI controllers. Other controller types are not supported. Also all the virtual disks must belong to a single datastore.

Additional prerequisites

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

- Using array-based replication
- Using Resiliency Platform Data Mover
- Using NetBackup
- For recovery on Amazon Web Services
- For recovery on vCloud

Using array-based replication

- The replicated storage must be provisioned to the ESX servers at each data center.
  See "Protecting VMware virtual machines using array-based replication - an overview" on page 42.
- Ensure that any file or device is not mounted on the virtual CD drive of the virtual machine.

Using Resiliency Platform Data Mover

- The Resiliency Platform Data Mover virtual appliances must be part of the same ESX clusters that host the protected virtual machines in the source and target data centers.
- Each host system in source and target data center must run ESXi version 6.0 U2 or later versions.
- Check the acceptance level on every host system in source and target data center. Ensure that the host system's acceptance level is not set to VMwareCertified. The allowed acceptance levels are VMwareAccepted, PartnerSupported, or CommunitySupported. By default, the ESX host is set to PartnerSupported.
- Virtual machine disks should be in the same datastore.
- Enough storage must be available on the recovery data center for provisioning the replicated virtual machines.
- Both the source and recovery (target) gateway must have external storage equivalent to 12GB for each virtual machine protected by the gateway pair. For

example, if a gateway pair supports 10 virtual machines, the source and recovery (target) gateway must each have 120GB of external storage.

- A maximum of 58 volumes or disks can be attached to the Replication Gateway. The total number of disks from the virtual machines in the resiliency group and the number of disks that are already attached to the gateway should not exceed more than 58. This is a VMware limitation.
- Configuring resiliency group for remote recovery using Data Mover fails if a virtual machine has snapshots. Hence remove all snapshots before proceeding with the operation. This is a VMware limitation.
- The VMware vCenter Server user account used to add the server to Resiliency Platform must have the appropriate privileges for Data Mover.

  See “VMware vCenter Server privileges required for implementing Resiliency Platform Data Mover” on page 55.
- Ensure that any file or device is not mounted on the virtual CD drive of the virtual machine.

See “Protecting VMware virtual machines using Resiliency Platform Data Mover - an overview” on page 57.

Using NetBackup

- Every asset selected to be in the resiliency group should have a NetBackup policy of type VMware.
- NetBackup master server on production and recovery data center should be running.
- The vCenters or ESX servers for the selected assets should be configured in NetBackup master server.
- There should be at least 1 backup schedule meeting the RPO defined in the service objective.
- Ensure that the datastore on the recovery data center has enough storage capacity.

See “Using NetBackup - an overview” on page 31.

For recovery on Amazon Web Services

- The Paravirtual (PV) Drivers for Windows must be installed. These drivers are located at:

  ```
  C:\ProgramData\Symantec\VRTSsfmh\spool\addons\store\
  VRTSitrptap-2.1.0.0\AWSPVDriver_7.4.3
  ```

After completing the **Prepare host for replication** task, follow the documentation of AWS for more information on driver installation.

- Enable the UUID for the virtual machines (disk.enableuuid=true).
- Ensure that the Replication Gateways have sufficient storage to handle the replication for the planned number of protected virtual machines.

  Both the on-premises gateway and the cloud gateway must have external storage equivalent to 12GB for each asset protected by the gateway pair. For example, if a gateway pair supports 10 virtual machines, the on-premises gateway and the cloud gateway must each have 120 GB of external storage.

  A maximum of 40 volumes or disks can be attached to the cloud Replication Gateway.
- On Windows hosts, initialize and reboot the disks that are attached to IDE controller.
- On Linux virtual machines, ensure that the NIC name in the file `ifcfg-Auto_ens192` or `ifcfg-Auto_eth0` matches with the actual NIC name on the system.
- A bucket must be created in S3 and must have a policy that assigns the following permissions:
  - s3:GetBucketLocation and s3:GetObject permissions on the bucket
  - ec2:ImportSnapshot, ec2:DescribeSnapshot, and ec2:CopySnapshot permissions on all the resources

  One role named ImportSnapshotRole must be created and the above policy is associated with this role. The service vmie.amazonaws.com should be able to assume this Role.

  For more information about the permissions required, refer to the AWS documentation.

  See "Sample policy statement for AWS" on page 172.

  See "Sample trust relationship for AWS" on page 173.

See "Using Resiliency Platform Data Mover for recovery to AWS- an overview" on page 33.

For recovery on vCloud

- Enable the UUID for the virtual machines (disk.enableuuid=true).
- Install one of the following:
  - VMware Tools
  - OpenVM Tools version 9.10 or later
  - OpenVM Tools version prior to 9.10 with DeployPkg package

- A maximum of 45 volumes or disks can be attached to the Replication Gateway. The total number of disks from the virtual machines in the resiliency group and the number of disks that are already attached to the gateway should not exceed more than 45.

See “Using Resiliency Platform Data Mover for recovery to vCloud- an overview” on page 37.

Limitations for virtual machine disaster recovery

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

Table 12-1

Limitations	Descriptions
Replication limitations	For more information on replication-based limitations of virtual machines, refer to the Hardware and Software Compatibility List (HSCL).
Limitations due to open-vm-tools on VMware virtual machines	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. To reconfigure the guest IP you need to uninstall open-vm-tools and install the latest version of VMware Tools. Note that the virtual machine should not have both VMware Tools and open-vm-tools installed on it. Virtual machines with physical RDM disks cannot be replicated using Resiliency Platform Data Mover.

Managing virtual machines for remote recovery (DR) using 3rd party replication technology

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

To manage virtual machines for remote recovery (DR) using 3rd party replication technology

1 Prerequisites

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

   See "Protecting VMware virtual machines using array-based replication - an overview" on page 42.

2 Navigate

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

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

3 Select the virtual machines:

   - Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.
   - Drag and drop virtual machines to **Selected Instances**.

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

5 The next page lists the service objectives that are available for the selected asset type. You can expand the service objective to view details. Select the service objective that provides disaster recovery operations.

6 Select the target (recovery) data center.

7 Select the target Cluster or ESXi host. This panel is displayed only if the replication technology is NetApp SnapMirror. But if the storage is mounted using NFS protocol on both the data centers, then this panel is not displayed.

   See "Target asset selection options" on page 91.

8. Select the non-replicating datastore on the source data center and map it to a datastore on the target data center.

   Click **Add selected datastore pair** to create a pair.

   If you are not using Raw Device Mapping (RDM) paths, then this panel is not displayed.

   This panel is displayed if the datastore having RDM path is not replicated. You then need to map the datastores on both the data centers.

9. Complete the network customization steps for the virtualization technology.

   See “Network customization options” on page 92.

10. Verify the summarized information and enter a name for the resiliency group.

When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the DR operation. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.

See “Viewing activities” on page 146.

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

See “Viewing resiliency group details” on page 80.

Target asset selection options

This panel is displayed only if the replication technology is NetApp SnapMirror or Resiliency Platform Data Mover.

Select the following in this panel:

- Select the target cluster.
  Each resiliency group must map to only one ESX cluster. The wizard validates which clusters on the target data center can meet the required number of virtual machines and disks.

- Review the hosts within the cluster.
  Once you select a cluster, the associated ESX hosts are displayed, and below them the datastores that are accessible from the ESX hosts.

- Select the datastore
  Storage that can be provisioned must be available in the selected datastore in order for the wizard to create the replicated disks on the target data center.
  Review the total disk size and compare this to the value in the **Free (GB)** column for the selected datastore.
  If there is insufficient memory, you can continue with the wizard and update the resources later.

Network customization options

Before you proceed with network customization, See “Prerequisites for network customization” on page 93.

You can do the following in this panel:

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

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

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

Since only IPv4 is supported, you may see a warning if there are IPv6 address: *Unable to apply IP customization some of the workloads.* Ignore this warning.

To customize the static IP of Windows guest virtual machines in the VMware environment, Resiliency Platform requires the user name and password to log on to the Windows virtual machines. This user name and password is to be specified under ‘Windows global user’. The user credentials can be Windows Active Directory user or Workgroup user. For Windows Active Directory user, the Active Directory should be common for both, the primary and the recovery data center.

If a Windows virtual machine is part of a Windows Active Directory, ensure that you log on to the virtual machine at-least once using the Active Directory credentials. This is applicable only if the recovery is on on-premises data center.

For more information refer to the *Deployment Guide*.

If the recovery data center is in cloud, then ensure that the IPs used for network customization are not already in use on the cloud.

If you choose to apply DNS customization, then you can add a host name to the virtual machine. Note that DNS customization is not supported for vCloud.

See “Prerequisites for network customization” on page 93.

Prerequisites for network customization

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

- IP, Gateway, Netmask, DNS, Domain Name, Mac address etc. information should present in the respective files of each network interface for which you want to customize the IP and DNS.
- If multiple network interfaces (NICs) are assigned to a virtual machine, then you need to apply IP customization to all the NICs.
- For virtual machines that are running on Linux ensure that NetworkManager and libvirtd service is in off state.
- The mac address configuration should be set as Manual/Static so that it does not change after the DR operation is performed. This is applicable only for recovery to premises data center.
- For Windows virtual machines, ensure that user access control (UAC) is disabled on the hosts.

Managing virtual machines for remote recovery (DR) using Resiliency Platform Data Mover

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

To manage virtual machines for remote recovery using Resiliency Platform Data Mover

1. Prerequisites

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

   See “Protecting VMware virtual machines using Resiliency Platform Data Mover - an overview” on page 57.

2. Navigate

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

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

3 Select the virtual machines:

- Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.
- Drag and drop virtual machines to **Selected Instances**.
  Ensure that each resiliency group is mapped to only one ESX cluster.

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

5 The next page lists the service objectives that are available for the selected asset type. You can expand the service objective to view details. Select the service objective that provides disaster recovery operations.

6 Select the target (recovery) data center.

See "Target asset selection options" on page 91.

7 Continue through the wizard to complete the configuration using Resiliency Platform Data Mover.

See "DR configuration options using Resiliency Platform Data Mover" on page 94.

8 Select the target EXSi host.

See "Target asset selection options" on page 91.

9 Complete the network customization steps for the virtualization technology.

See "Network customization options" on page 92.

10 Verify the summarized information and enter a name for the resiliency group.

11 When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the DR operation. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.

See "Viewing activities" on page 146.

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

See "Viewing resiliency group details" on page 80.

DR configuration options using Resiliency Platform Data Mover

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

Information that is requested and displayed on the panel depends on the remote data center selection.

Table 12-2 Options for Resiliency Platform Data Mover

Wizard steps and options	Description
Select Replication Gateway Pair	Each gateway pairs provides information on the available assets and disks. If the gateways in the pair have different available capacities, the smaller capacity assets and disks are shown. Select a gateway pair to use for this resiliency group. Ensure that the available assets and disks are sufficient for the number of assets and disks listed as the **Assets requirement**. If remote recovery is in Amazon Web Services (AWS), then the maximum disks that can be attached to the gateway is 40. If remote recovery is in vCloud, then the maximum disks that can be attached to the gateway is 45.
Confirm Data Mover Details	Verify the replication gateway pair selection and the assets before you continue with the wizard.
Enable reverse replication after migrate operation	This option is displayed only if the recovery data center is in cloud. This option starts the reverse replication of data after successfully migrating the virtual machines to the recovery data center. If the check box is selected then Resync is not required, else you need to run the Resync operation before migrating the virtual machines back to the production data center. Do not run the Resync operation while the reverse replication is in progress.

Managing virtual machines for remote recovery (DR) in Amazon Web Services

Using the Resiliency Platform console, you can organize virtual machines into a resiliency group, apply the remote recovery for hosts service objective, and configure them for remote recovery in Amazon Web Services (AWS).

The wizard prompts for the inputs that are needed for the selected service objective and replication technology.

To manage virtual machines for remote recovery in AWS

1 Prerequisites

2 Navigate

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

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

3 Select the virtual machines:

- Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.
- Drag and drop virtual machines to **Selected Instances**.
  Ensure that each resiliency group is mapped to only one ESX cluster.

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

5 The next page lists the service objectives that are available for the selected asset type. You can expand the service objective to view details. Select the service objective that provides disaster recovery operations.

6 Select the target (recovery) data center.

7 Continue through the wizard to configure Resiliency Platform Data Mover for replication.

See “DR configuration options using Resiliency Platform Data Mover” on page 94.

8 Select the target volume type for each disk. Enter the IOPS required if the volume type is Provisioned IOPS SSD.

Refer to AWS documentation for more information on IOPS permitted for specific volume type and size.

9 Complete the customization for AWS.

See "AWS Customization options panel" on page 97.

10 Complete the network customization steps for the virtualization technology.

See "Network customization options" on page 92.

11 Verify the summarized information and enter a name for the resiliency group.

12 When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the DR operation. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.

See "Viewing activities" on page 146.

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

See "Viewing resiliency group details" on page 80.

If the operation fails and you want to delete the resiliency groups, and while deleting the resiliency group if you face any issues, you need to perform manual cleanup steps.See "Troubleshooting delete resiliency group operation" on page 168.

AWS Customization options panel

This wizard panel is displayed when you are configuring the assets for remote recovery in Amazon Web Services (AWS).

The following table summarizes the information that you must supply or verify in the wizard to complete the customization.

Table 12-3 Options for AWS customization

Wizard steps and options	Description
Select Attributes	Review the availability zone for the selected virtual machines. Select a security group and the instance type. You can choose multiple security groups for a virtual machine. Ensure that the security group is in the same VPC as the subnet that is selected during network mapping. After selecting the security group and the instance type you can choose to apply the selection to each virtual machine or to all.

Table 12-3 Options for AWS customization *(continued)*

Wizard steps and options	Description
Network Details	Displays the source and the target network mapping details that are applicable for the selected resiliency groups. If network mapping is not done, then the page is empty.

Managing virtual machines for remote recovery (DR) in vCloud

Using the Resiliency Platform console, you can organize virtual machines into a resiliency group, apply the remote recovery for hosts service objective, and configure them for remote recovery in vCloud.

The wizard prompts for the inputs that are needed for the selected service objective and replication technology.

To manage virtual machines for remote recovery in vCloud

1 Prerequisites

See “Prerequisites for configuring VMware virtual machines for disaster recovery” on page 85.

2 Navigate

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

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

3 Select the virtual machines:

- Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.
- Drag and drop virtual machines to **Selected Instances**.
  Ensure that each resiliency group is mapped to only one ESX cluster.

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

5 The next page lists the service objectives that are available for the selected asset type. You can expand the service objective to view details. Select the service objective that provides disaster recovery operations.

6 Select the target (recovery) data center.

7 Continue through the wizard to configure Resiliency Platform Data Mover for replication.

See “DR configuration options using Resiliency Platform Data Mover” on page 94.

8 In the **vCloud Configuration** panel, select the storage profile on target data center for provisioning storage.

You can apply the selection to all virtual machines or select a storage profile for each virtual machine.

9 In the **Customization** panel review the target data center details and click **Next**.

10 Complete the network customization steps for the virtualization technology.

See “Network customization options” on page 92.

11 Verify the summarized information and enter a name for the resiliency group.

12 When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the DR operation. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.

See “Viewing activities” on page 146.

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

See “Viewing resiliency group details” on page 80.

Managing VMware virtual machines for remote recovery using NetBackup images

When you want to protect your assets that are being backed up using NetBackup, you need to organize them into a resiliency group and apply the service objective where the data availability mode is Copy.

To manage virtual machines for remote recovery using NetBackup images

1 Prerequisites

- Ensure that you have completed the NetBackup configuration tasks and the Local and remote service objective, with data availability as Copy, is activated.

  See “Using NetBackup - an overview” on page 31.

- Ensure that the virtual machine configuration prerequisites are met.
  See “Prerequisites for configuring VMware virtual machines for disaster recovery” on page 85.

2. Navigate

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

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

3. Select the virtual machines:
   - Select **Host** as the asset type, select the data center, and select other filters as needed to display a list of virtual machines.
   - Select the virtual machines that are backed up using NetBackup. You can identify the assets based on the NetBackup policy name and the technology that is displayed in the **Data availability** column.
     Drag and drop the required virtual machines to **Selected Instances**.
4. The next page displays the environment for the selected assets on both the data centers.
5. The next page lists the service objectives that are available for the selected assets. You can expand the service objective to view details. Select the appropriate service objective.
6. Select the target (recovery) data center.
7. Do the following in the NetBackup Configuration panel and click **Next**.
   - Review the information about common policies across the selected virtual machines and the master server on which the data is to be restored. Review the target master servers selection, the schedule, and policy details.
   - Select the target ESX server and datastore for all the selected assets or for each asset. You can choose the ESX server and the datastore on which the virtual machine is to be restored remotely.
   - Review the summary.
8. Complete the network customization steps for the virtualization technology.

   See “Network customization options” on page 92.
9. Verify the summarized information and enter a name for the resiliency group.

When you finish the wizard steps, Resiliency Platform invokes a workflow which initializes the virtual machines for the DR operation. You can view the progress or ensure that this operation is successfully completed on the **Activities** page.

See “Viewing activities” on page 146.

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

See “Viewing resiliency group details” on page 80.

Note that after configuring the resiliency group, it might take some time for the backup images to reflect in the Restore Wizard. Also note that all the available images from NetBackup master server are not discovered. The discovery of images is on the selected RPO which is defined in the service objective. If the RPO is “X”, at the time of creation of resiliency group, then 2X or past one month backup images, whichever is the maximum duration are discovered. Example if the RPO is 2 months, then the images for last 4 months since the time of creation of resiliency group are discovered. If the RPO is 4 hours then images for last one month are discovered.

Verifying the replication status for Veritas Resiliency Platform Data Mover

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

See “Viewing resiliency group details” on page 80.

To verify the replication status for Veritas Resiliency Platform Data Mover

1 Navigate

Assets (navigation panel)

Resiliency Groups

2 Double-click the resiliency group for which DR using Resiliency Platform Data Mover is already configured.

3 View the replication state at **Replication** > **State**.

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

Some possible states:

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

Table 12-4 Data Mover replication states

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

Table 12-5 Replication states

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

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

- Migrate operation: Is allowed when replication state is **Connected** and data state is **Consistent** or **Inconsistent Good**.

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

Section 4

Managing disaster recovery

- Chapter 13. Rehearsing DR operations to ensure DR readiness
- Chapter 14. Performing disaster recovery operations

Chapter 13

Rehearsing DR operations to ensure DR readiness

This chapter includes the following topics:

- About ensuring the disaster recovery readiness of your assets
- Rehearse operations - array-based replication
- Rehearse operations - Resiliency Platform Data Mover
- Prerequisites for rehearsal operation
- Performing the rehearsal operation
- Performing the rehearsal operation using NetBackup images
- Performing cleanup rehearsal

About ensuring the disaster recovery readiness of your assets

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

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

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

Rehearse operations - array-based replication

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

Rehearse operations with EMC SRDF-based replication

Rehearse operations with NetApp SnapMirror based replication

Rehearse operations with EMC SRDF-based replication

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

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

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

Rehearse operations with NetApp SnapMirror based replication

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

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

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

Rehearse operations - Resiliency Platform Data Mover

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

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

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

See “Viewing and configuring network settings for a data center” on page 60.

See “Setting up network mapping between production and recovery data centers” on page 64.

Prerequisites for rehearsal operation

Before you run the rehearsal operation for a resiliency group, ensure that you have met the following prerequisites:

- For VMware virtual machines, ensure that the datastores have enough free space for the swap files for the on-premises virtual machines and the virtual machines created by the rehearsal operation on the recovery data center. The size of the swap files is same as that of the virtual machine memory size.
- For VMware virtual machines, ensure that the mapping of all the required port groups (VLANs) across the data centers is complete.
  See “Setting up network mapping between production and recovery data centers” on page 64.
- Each type of replication has prerequisites and limitations for the rehearsal operation.
  See “Rehearse operations - array-based replication” on page 106.
  See “Rehearse operations - Resiliency Platform Data Mover” on page 107.
- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs.
- If the recovery data center is in AWS, then configure a rehearsal subnet in the cloud. The rehearsal and production subnet should be in the same VPC.

Performing the rehearsal operation

Use the **Rehearsal** option on the Resiliency Platform console to ensure the disaster recovery readiness of the assets in your protected resiliency groups.

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

For recovery on AWS cloud:

The time taken to complete the Rehearsal operation depends on the size and the number of volumes. If the recovery data center is in AWS cloud, then to reduce the time taken to complete the snapshot creation task during Rehearsal, you may take a snapshot of the volumes manually before running the Rehearsal operation. Before taking a snapshot, ensure that the replication state is Consistent. Since, in AWS the subsequent snapshots are only incremental, the time taken to create snapshots during Rehearsal is significantly reduced. Which reduces the overall time taken to complete the operation.

To perform the rehearsal operation

1 Prerequisites

See “Prerequisites for rehearsal operation” on page 108.

2 Navigate

Assets (navigation pane)

Resiliency Groups

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

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

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

See “Performing cleanup rehearsal ” on page 110.

Performing the rehearsal operation using NetBackup images

Configure VMware virtual machines for protection using a service objective where the data availability mode is Copy. Use the **Rehearsal** option on the Resiliency Platform console to ensure the disaster recovery readiness of the assets in your protected resiliency groups.

Then during the rehearsal operation virtual machines are created on the recovery data center with the selected backup image. Note that the virtual machines that are created on the recovery data center have a different name and are in a different folder than the virtual machines on the production data center.

To perform the rehearsal operation using NetBackup images

1 Navigate

Assets (navigation pane)

Resiliency Groups

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

3 Select the target data center and the type of image to be used. You can choose the latest image or select one from the required time range.

4 Click **Next** and **Finish** to submit the wizard.

Performing cleanup rehearsal

After you have performed the rehearsal operation successfully to verify the ability of your configured resiliency group to fail over on to the disaster recovery data center, you can use the cleanup rehearsal operation to clean up the rehearsal virtual machines or applications in the resiliency group. All temporary objects created during the rehearsal operation are now deleted.

Using NetBackup

When your assets are configured for remote recovery using a service objective where the data availability mode is Copy, then during the rehearsal operation virtual machines are created on the recovery data center with the selected backup image. These virtual machines and the data are deleted during the cleanup operation.

To perform cleanup rehearsal

1 Navigate

Assets (navigation pane)

Resiliency Groups

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

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

If the replication technology used is 3PAR Remote Copy, then refresh the 3PAR enclosure after successfully completing the rehearsal cleanup operation.

See “Performing the rehearsal operation” on page 108.

Chapter 14

Performing disaster recovery operations

This chapter includes the following topics:

- How Resiliency Platform Data Mover handles DR operations
- Migrating a resiliency group of virtual machines
- Taking over a resiliency group of virtual machines
- Performing the resync operation
- Restoring data using NetBackup
- Clearing outage

How Resiliency Platform Data Mover handles DR operations

Migration

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

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

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

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

Take over

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

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

Since it is an unplanned event, the data available at the recovery data center may not be up-to-date. You need to evaluate the tolerable limit of data loss. If the available data is within the acceptable limits, perform the takeover operation to bring up the source workloads in the recovery data center. The takeover operation provisions and brings up the virtual machines at the recovery data center using the latest data on the recovery data center storage.

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

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

Migrating workloads back to the source data center after a takeover

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

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

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

Migrating a resiliency group of virtual machines

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

If you perform the takeover operation, then you must perform the resync operation before you migrate back to the production data center.

If the **Enable reverse replication** option is not selected while configuring for remote recovery, then you need to run the Resync operation before migrating the virtual machines back to the production data center.

See “Performing the resync operation” on page 115.

If the recovery data center is AWS cloud, then before you migrate from the cloud data center to the on-premises data center, you need to reboot and then refresh the virtual machine in the cloud.

To migrate virtual machines

1
   - Ensure that the replication is in Active state and the data is consistent.
   - It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs.
   - For VMware virtual machines, ensure that the network mapping of all the required port groups (VLANs), or subnets across the data centers is complete.
     See “Setting up network mapping between production and recovery data centers” on page 64.
   - If the recovery data center is in AWS, then ensure that the network mapping of all the required subnets across the data centers is complete.

2 Navigate

   Assets (navigation pane)

   Resiliency Groups

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

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

If the Migrate operation fails, check **Recent Activities** to know the reason and fix it. You can then launch the **Retry** operation. The **Retry** operation restarts the migrate workflow, it skips the steps that were successfully completed and retries those that had failed.

Do not restart the workflow service while any workflow is in running state, otherwise the **Retry** operation may not work as expected.

Taking over a resiliency group of virtual machines

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

Perform the resync operation after successful completion of takeover operation.

If the recovery data center is in cloud, then takeover operation from cloud data center to production (on-premises) data center is not supported.

To perform takeover operation on virtual machines

1 Prerequisites

- It is recommended to stop or disable NetworkManager on RHEL hosts having multiple NICs.
- For VMware virtual machines, ensure that the network mapping of all the required port groups (VLANs), or subnets across the data centers is complete.
  See "Setting up network mapping between production and recovery data centers" on page 64.

- If the recovery data center is in AWS, then ensure that the network mapping of all the required subnets between the production and recovery data center is complete.

2 Navigate

Assets (navigation pane)

Resiliency Groups

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

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

If the Takeover operation fails, check **Recent Activities** to know the reason and fix it. You can then launch the **Retry** operation. The **Retry** operation restarts the migrate workflow, it skips the steps that were successfully completed and retries those that had failed.

Do not restart the workflow service while any workflow is in running state, otherwise the **Retry** operation may not work as expected.

Performing the resync operation

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

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

Use the Resync operation on the Resiliency Platform console to automate these steps for the required resiliency groups. This operation cleans up the residue which includes stopping applications and virtual machines, unregistering virtual machines, unmounting file systems, datastores, etc. If the recovery data center is Amazon Web Services, then the virtual machines are not unregistered.

If the recovery data center is in cloud, and if while configuring the resiliency group for remote recovery, the **Enable reverse replication** option was selected, then do not run the Resync operation till the replication is complete and the state is IN Sync.

Performing the resync operation

1 Prerequisites

If the recovery data center is on-premises, then restart the ESX servers on production data center before performing resync operation. Restarting the ESX servers ensures that all stale references to virtual machines, disks, or datastores are released so that resync can work properly.

2 Navigate

Assets (navigation pane)

Resiliency Groups

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

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

If the Resync operation fails, check **Recent Activities** to know the reason and fix it. You can then launch the **Retry** operation. The **Retry** operation restarts the migrate workflow, it skips the steps that were successfully completed and retries those that had failed.

Do not restart the workflow service while any workflow is in running state, otherwise the **Retry** operation may not work as expected.

Restoring data using NetBackup

Using Resiliency Platform console, you can recover VMware virtual machines that are protected by NetBackup and managed using a service objective where the data availability mode is Copy.

You can choose from the most recent successful backup image or another available backup image.

You can restore data on VMware virtual machines in the same data center or in a remote data center.

A virtual machine is restored using the selected backup image at selected data center, which could be the same data center or a remote data center. Before that, the virtual machine whose data is being restore is stopped and unregistered. Virtual machines are unregistered only if the restore is on a remote data center.

The network settings of the virtual machine are configured as per the subnet mapping, switch or port mapping as specified during the configure for remote recovery operation. Finally the virtual machine is up and running at the recovery/selected data center.

The state of the virtual machine is **Deleted/Unavailable**.

If a virtual machine is deleted, you can choose the latest backup copy or any previous backup to restore the virtual machine on the required data center.

After configuring a resiliency group for recovery using a service objective where the data availability mode is Copy, it might take some time for the backup images to reflect in the Restore wizard.

When you choose the Local recovery of hosts service objective, then during the restore operation remote data center is not displayed. Using this service objective you can choose the virtual machines whose data you want to restore irrespective of the number of virtual machine in the resiliency group.

For Local and remote recovery of hosts service objective, if virtual machines are registered in the target data center, then restore is of the type local. If virtual machines are not registered in the target data center, then the restore is of the type remote.

After the restore operation is successfully completed, you may refresh the NetBackup master server to discover new backup policies. Or you may wait for the scheduled discovery to complete.

Consider the following before running the restore operation:

- If a virtual machine is already present on the recovery site and is online, then the data on the virtual machine is not overwritten during the restore operation.
- If the virtual machine is registered but in offline state, then the virtual machine is deleted. A new virtual machine is created and the backup image is restored on that machine.
- If the virtual machine is present in the inventory but not registered, then the virtual machine is deleted. A new virtual machine is created and the backup image is restored on that machine.
- If the virtual machine is not registered and is not present in the inventory, then a new virtual machine is created and the backup image is restored on that virtual machine.

A successfully restored image is tagged. Hover the mouse on the tag to view additional information such as the NetBackup master server name, the image type (full or incremental), and the last test date and time.

To restore virtual machine data at the recovery data center

1 Prerequisites

- The resiliency group that contains the VMware virtual machines must be protected using the service objective activated for copy based recovery.
- The IMS, ESX server, and the vCenter on the recovery data center must be connected to the NetBackup master server.
- Ensure that the state of NetBackup master server is Connected which means that the notification channel is working.

2. Navigate

   Assets > **Resiliency Groups** tab

3. Go to the details page for the resiliency group with the virtual machines that you want to recover and select **Restore**.
4. In the **Restore Resiliency Group** panel do the following and click **Next**.
   - Review the target data center.
   - Select **Confirm outage of assets for resiliency group** to confirm that there is an outage on the selected data center.

     This is not applicable if the restore is of the type local.
   - Choose the backup image to be used. You can choose from the latest image or any other available backup from the specified time range.

     To choose from a specific time range, select the date and the time range. Then select the hours or minutes since the start time. Click **Search**.
5. Click **Finish**.

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

On successful completion of the restore operation, verify that the selected virtual machines are online.

Clearing outage

While performing the restore operation if an outage was declared on the production data center, then you need to clear the outage. Only after you have cleared the outage, can you restore the backup images from the remote data center to the production data center.

To clear an outage

1. Navigate

   Assets > **Resiliency Groups** tab

2. Go to the details page for the resiliency group and select **Clear outage**.
3. In the **Clear outage** panel review the information and click **Next**.
4. Click **Finish**.

Managing resiliency plans

This chapter includes the following topics:

- About resiliency plans
- Creating a new resiliency plan template
- Editing a resiliency plan template
- Deleting a resiliency plan template
- Viewing a resiliency plan template
- Creating a new resiliency plan
- Editing a resiliency plan
- Deleting a resiliency plan
- Executing a resiliency plan
- Viewing a resiliency plan
- Creating a schedule for a resiliency plan
- Editing a schedule for a resiliency plan
- Deleting a schedule for a resiliency plan
- Viewing a schedule for a resiliency plan

About resiliency plans

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

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

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

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

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

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

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

See “Creating a new resiliency plan template” on page 121.

See “Creating a new resiliency plan” on page 126.

Creating a new resiliency plan template

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

- Start and stop a resiliency group.
- Rehearsal and cleanup rehearsal of a resiliency group.
- Migrate and takeover a resiliency group.
- Manual task
  See “About manual task” on page 122.
- Run a custom script
  See “About custom script” on page 123.

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

   Automation Plans (menu bar) > **Resiliency Plans** 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 120.

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 121.

  See "Creating a new resiliency plan" on page 126.

2 Drag and drop **Manual Task** into the canvas. Click the pencil icon in the action box to add the task details.

3 Provide a name for the manual task.

4 Describe the reason why you want to add this manual task to the resilient plan.

5 Select your choice for resuming the process manually or automatically. If you select the option for automatically resuming the process after a timeout, enter the duration of timeout in minutes. Click **Save**.

About custom script

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

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

The custom script can be in any format that can be directly executed on a shell on the target host. For the Linux hosts, it may be an executable or a script that specifies the interpreter on the hashbang 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 an EXE. 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 121.

   See "Creating a new resiliency plan" on page 126.

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

   Automation Plans (menu bar) > **Resiliency Plans** 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 121.

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

   Automation Plans (menu bar) > **Resiliency Plans** 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 121.

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

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

2 In the **Templates** list, do one of the following:

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

3 You can now view the details of the resiliency plan template.

Creating a new resiliency plan

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

- Start and stop a resiliency group.
- Rehearsal and cleanup rehearsal of a resiliency group.
- Migrate and takeover a resiliency group.
- Manual task
  See "About manual task" on page 122.
- Run a custom script
  See "About custom script" on page 123.

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

To create a new resiliency plan using blank template

1. Navigate

   Automation Plans (menu bar) > **Resiliency Plans** 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 120.

See “Deleting a resiliency plan” on page 128.

See “Executing a resiliency plan” on page 128.

Editing a resiliency plan

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

To edit a resiliency plan

1 Navigate

Automation Plans (menu bar) > **Resiliency Plans** 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 126.

Deleting a resiliency plan

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

To delete a resiliency plan

1 Navigate

Automation Plans (menu bar) > **Resiliency Plans** 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 126.

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

   Automation Plans (menu bar) > **Resiliency Plans** 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 126.

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 128.

See "Deleting a resiliency plan" on page 128.

To view a resiliency plan

1. Navigate

   Automation Plans (menu bar) > **Resiliency Plans** 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

Automation Plans (menu bar) > **Resiliency Plans** 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

Automation Plans (menu bar) > **Resiliency Plans** 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

   Automation Plans (menu bar) > **Resiliency Plans** 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 130.

See “Deleting a schedule for a resiliency plan” on page 130.

To view a schedule for a resiliency plan

1. Navigate

   Automation Plans (menu bar) > **Resiliency Plans** 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.

Chapter 16

Monitoring risks, reports, and activities

This chapter includes the following topics:

- About the Resiliency Platform Dashboard
- Understanding asset types
- Displaying an overview of your assets
- About risk insight
- Viewing reports
- Managing activities

About the Resiliency Platform Dashboard

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

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

The Dashboard has the following areas:

Global View	A world map that identifies the data centers that contain Resiliency Platform managed assets. Lines between data centers indicate that replication takes place between the locations. Mouse over an icon for basic Resiliency Platform configuration and asset configuration information for that data center. Click **More** for detailed information and recent activity.
Resiliency Groups and **Virtual Business Services** summaries	The upper right section of the dashboard displays total number of resiliency groups and virtual business services in the resiliency domain, as well as those at risk and normal. Click a square in either the **Resiliency Groups** or **Virtual Business Services** summary to display a tab of detailed information. The **Activity Summary** provides details of the DR activities such as average time taken, failed and successful runs.
Virtual Machines by Platform and OS	Displays a summary of virtual machines in all data centers or information on a single data center. Use the drop-down list to filter your results. The summary lists the virtual machine types by percentage and the platform types by number.
Risks Summary	Displays a summary of errors and warning in all data centers. Click **View Details** to view additional information.
Application environment	Displays the number of applications and the application types. The chart shows the number of applications that are managed by InfoScale and those that are not managed by InfoScale.
Applications by Type	Displays a summary of application types in all data centers or in a single data center. Use the drop-down list to filter your results.
Top Resiliency Groups by Replication Lag	Ranks the resiliency groups according to how long it takes the recovery data center to be in sync with the active data center.

By Service Objective	Displays the percentage of virtual machines and applications that are unprotected or unmanaged. Use the drop-down list to filter your results.

See “Displaying resiliency group information and status” on page 77.

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

Displaying an overview of your assets

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

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

🗔

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

- Unprotected resiliency groups, are groups under Resiliency Platform control, but that do not have disaster recovery configured.

  See "Managing virtual machines for basic monitoring" on page 75.

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

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

For VBSs, the screen displays the following:

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

About risk insight

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

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

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

Risk insight generates two types of reports:

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

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

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

Table 16-1

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 16-2

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.

See “Displaying risk information” on page 136.

See “Predefined risks in Resiliency Platform” on page 137.

See “Viewing the current risk report” on page 143.

See “Viewing the historical risk report” on page 144.

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 16-3 Ways to display risks

To display ...	Do the following:
A complete list of risks across the resiliency domain	**1** On the menu bar, select ⊞ **More Views** > **Risks** **2** On the **Risk** page, double-click a risk in the table to display detailed information.
Risks that are associated with a specific resiliency group or virtual business service	**1** On the navigation pane, select 🗔 (Assets) and the tab for either **Resiliency Groups** or **Virtual Business Services**. **2** On the tab, double-click a resiliency group or virtual business service to display detailed information. **3** On the details page, note any risks that are listed in the **At Risk** area, and double-click the risk for details.

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

Predefined risks in Resiliency Platform

Table 16-4 lists the predefined risks available in Resiliency Platform. These risks are reflected in the current risk report and the historical risk report.

Table 16-4 Predefined risks

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Veritas Infoscale Operations Manager disconnected	Checks for Veritas Infoscale Operations Manager to Resiliency Manager connection state	1 minute	Error	All operations	Check Veritas Infoscale Operations Manager reachability Try to reconnect Veritas Infoscale Operations Manager
vCenter Password Incorrect	Checks if vCenter password is incorrect	5 minutes	Error	■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations	In case of a password change, resolve the password issue and refresh the vCenter configuration
VM tools not installed	Checks if VM Tools are not Installed. It may affect IP Customization and VM Shutdown.	Real time, when resiliency group is created	Error	■ Migrate ■ Stop	■ In case of VMWare, install VMWare Tools ■ In case of Hyper-V, install Hyper-V Integration Tools
Snapshot removed from Virtual Machine	Checks if snapshot has been removed from virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Edit the resiliency group to refresh configuration
Snapshot reverted on Virtual Machine	Checks if snapshot has been reverted on virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Remove and re-add the virtual machine to the Resiliency group by editing Resiliency group

Table 16-4 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Data Mover Daemon Crash	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX.	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Snapshot created on Virtual Machine	Checks if a snapshot has been created on Virtual machine.	5 minutes	Error	Resiliency Platform Data Mover replication	Edit the resiliency group to refresh configuration
DataMover virtual machine in noop mode	Checks if VM Data Mover filter is not able to connect to its counterpart in ESX.	5 minutes	Error	Resiliency Platform Data Mover replication	In order to continue the replication, you can move (VMotion) the VM to a different ESX node in the cluster and either troubleshoot the issue with this ESX node or raise a support case with Veritas
Resiliency group configuration drift	Checks if disk configuration of any of the assets in the resiliency group has changed.	30 minutes	Error	■ Migrate ■ Resync	Edit the resiliency group to first remove the impacted virtual machine from the resiliency group and then add it back to the resiliency group.
Global user deleted	Checks if there are no global users. In this case, the user will not be able to customize the IP for Windows machines in VMware environment.	Real time	Warning	■ Migrate ■ Takeover	Edit the resiliency group or add a Global user

Table 16-4 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Missing heartbeat from Resiliency Manager	Checks for heartbeat failure from a Resiliency Manager.	5 minutes	Error	All	Fix the Resiliency Manager connectivity issue
Infrastructure Management Server disconnected	Check for Infrastructure Management Server(IMS) to Resiliency Manager(RM) connection state.	1 minute	Error	All	Check IMS reachability Try to reconnect IMS
Storage Discovery Host down	Checks if the discovery daemon is down on the storage discovery host	15 minutes	Error	Migrate	Resolve the discovery daemon issue
DNS removed	Checks if DNS is removed from the resiliency group where DNS customization is enabled	real time	Warning	■ Migrate ■ Takeover	Edit the Resiliency Group and disable DNS customization
IOTap driver not configured	Checks if the IOTap driver is not configured	2 hours	Error	None	Configure the IOTap driver This risk is removed when the workload is configured for disaster recovery
VMware Discovery Host Down	Checks if the discovery daemon is down on the VMware Discovery Host	15 minutes	Error	Migrate	Resolve the discovery daemon issue
VM restart is pending	Checks if the VM has not been restarted after add host operation	2 hours	Error	Configure DR	Restart the VM after add host operation
New VM added to replication storage	Checks if a virtual machine that is added to a Veritas Replication Set on a primary site, is not a part of the resiliency group.	5 minutes	Error	■ Migrate ■ Takeover ■ Rehearsal	Add the virtual machine to the resiliency group.

Table 16-4 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Replication lag exceeding RPO	Checks if the replication lag exceeds the thresholds defined for the resiliency group. This risk affects the SLA for the services running on your production data center.	5 minutes	Warning	■ Migrate ■ Takeover	Check if the replication lag exceeds the RPO that is defined in the Service Objective
Replication state broken/critical	Checks if the replication is not working or is in a critical condition for each resiliency group.	5 minutes	Error	■ Migrate ■ Takeover	Contact the enclosure vendor.
Remote mount point already mounted	Checks if the mount point is not available for mounting on target site for any of the following reasons: ■ Mount point is already mounted. ■ Mount point is being used by other assets.	■ Native (ext3, ext4,NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours	Warning	■ Migrate ■ Takeover	Unmount the mount point that is already mounted or is being used by other assets.
Disk utilization critical	Checks if at least 80% of the disk capacity is being utilized. The risk is generated for all the resiliency groups associated with that particular file system.	■ Native (ext3, ext4,NTFS): 30 minutes ■ Virtualization (VMFS, NFS): 6 hours	Warning	■ Migrate ■ Takeover ■ Rehearsal	Delete or move some files or uninstall some non-critical applications to free up some disk space.
ESX not reachable	Checks if the ESX server is in a disconnected state.	5 minutes	Error	■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations	Resolve the ESX server connection issue.

Table 16-4 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
vCenter Server not reachable	Checks if the virtualization server is unreachable or if the password for the virtualization server has changed.	5 minutes	Error	■ On primary site: start or stop operations ■ On secondary site: migrate or takeover operations	Resolve the virtualization server connection issue. In case of a password change, resolve the password issue.
Insufficient compute resources on failover target	Checks if there are insufficient CPU resources on failover target in a virtual environment.	6 hours	Warning	■ Migrate ■ Takeover	Reduce the number of CPUs assigned to the virtual machines on the primary site to match the available CPU resources on failover target.
Host not added on recovery data center	Checks if the host is not added to the IMS on the recovery data center.	30 minutes	Error	Migrate	Check the following and fix: ■ Host is up on recovery data center. ■ Host is accessible from recovery datacenter IMS. ■ Time is synchronized between host and recovery datacenter IMS.
NetBackup Notification channel disconnected	Checks for NetBackup Notification channel connection state	5 minutes	Error	Restore	Check if the NetBackup Notification channel is added to the NetBackup master server.

Table 16-4 Predefined risks *(continued)*

Risks	Description	Risk detection time	Risk type	Affected operation	Fix if violated
Backup image violates the defined RPO	Checks if the backup image violates the defined RPO	30 minutes	Warning	No operation	■ Check the connection state of NetBackup Notification channel ■ Check for issues due to which backup images are not available
NetBackup master server disconnected	Checks if NetBackup master server is disconnected or not reachable	5 minutes	Error	Restore	Check if IMS is added as an additional server to the NetBackup master server
Assets do not have copy policy	Checks if the assets do not have a copy policy	3 hours	Warning	No operation	Set up copy policy and then refresh the NetBackup master server
Target replication is not configured	Checks if the target replication is not configured	3 hours	Warning	No operation	Configure target replication and then refresh the NetBackup master server
Disabled NetBackup Policy	NetBackup policy associated with the virtual machine is disabled	3 hours	Warning	No operation	Fix the disabled policy

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.

Viewing reports

Veritas Resiliency Platform provides a console for viewing the following reports:

Resiliency Groups and VBS Summary

Provides details about the resiliency groups and VBSs in the data centers across all sites.

VM Inventory

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.

The **Details** table provides additional information for each virtual machine.

Virtual Infrastructure Inventory

Provides information about the virtual infrastructure inventory across data centers. A pie chart shows the platform and virtualization technology distribution of the virtual servers across all data centers.

The **Details** table provides additional information for each virtual server.

Activity Distribution History

Provides information about tasks, such as migrate, takeover, rehearse, start, and stop, performed for a specified duration.

Recovery Activity History by RG

Provides historical information about recovery tasks, such as migrate, takeover, rehearse, and restore for each resiliency group.

Recovery Activity History by VBS

Provides historical information about recovery tasks, such as migrate, takeover, rehearse, and restore for each VBS.

Metering

Provides details of the virtualization servers that are protected for disaster recovery.

You can view the total number of servers that are protected for disaster recovery. For these servers you can view the total memory, processor cores, and the total storage.

VBS RPO	Provides Recovery Point Objective (RPO) details for all the virtual business services (VBS) in the resiliency domain. The bar chart provides information on the top VBS with maximum RPO lag. You can view the lag in the last replication and the replication date for all the VBS in the table.

To view a report

1. 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*.

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 146.

See “Aborting a running activity” on page 147.

Viewing activities

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

Click on a currently running task, to view the details in a graphical representation. The steps that are completed are shown in green color along with the success icon.

The ongoing steps are in blue color with the loader image, and the future steps are in gray. Expand **Execution Details** to view all the sub-tasks that comprise the task.

To view activities

1 Navigate

⊞ **Activities** (menu bar).

2 Choose either of the following:

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

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

See “Aborting a running activity” on page 147.

Aborting a running activity

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

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

To abort an activity

1 Navigate

Do one of the following:

⊞ **Activities**. Skip to 2

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

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

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

See “Viewing activities” on page 146.

Chapter 17

Managing evacuation plans

This chapter includes the following topics:

- About evacuation plan
- Generating an evacuation plan
- Regenerating an evacuation plan
- Performing evacuation
- Performing rehearse evacuation
- Performing cleanup evacuation rehearsal

About evacuation plan

An evacuation plan lets you evacuate all the assets from the production data center to the recovery data center with a single click operation.

Using the evacuation plan template you can define the sequence in which the virtual business services (VBS) should be migrated from the production data center to the recovery data center. Resiliency groups that do not belong to any VBSs, are appended at the end of the evacuation plan workflow after the VBS.

You can create an evacuation plan using only resiliency groups also. Having a VBS is not compulsory.

An evacuation plan has Priorities. You can add the VBSs to different priority levels. Ordering of resiliency groups is done by the Resiliency Platform.

If an asset within a VBS or a resiliency group fails to recover, the evacuation plan skips the asset and continues the process for the remaining assets. To do this you need to select the **Continue on failures** check box while creating the evacuation plan.

If the check box is not selected the evacuation plan stops, enabling you to fix the problem, and proceed ahead. If you choose to restart the workflow then the already executed steps are re-executed with the same results.

Only users with **Manage Evacuation Plans** permission can create and run the evacuation plans.

- VBS or resiliency group that belong to the evacuation plan must be configured for disaster recovery.
- VBS can contain resiliency groups some of which are configured for disaster recovery and some using the service objective with data availability as Copy.
- Resiliency group must belong to only one VBS.

When you generate a plan, an appropriate warning is shown listing the assets that are excluded from the plan.

On completing the evacuation plan, you can perform the following operations:

- Evacuate
- Rehearse evacuation
- Cleanup evacuation rehearsal
- Regenerate

An alert is raised and you need to perform the **Regenerate evacuation plan** operation in the following scenarios:

- VBSs are added, modified, or deleted.
- Resiliency groups are added and configured for disaster recovery.
- Resiliency groups which were configured for disaster recovery are deleted.
- Existing resiliency group is configured for disaster recovery.

No action is required in the following scenarios:

- Resiliency groups are modified.
- Resiliency groups which are not configured for disaster recovery are deleted.

When you run the **Evacuate**, **Rehearse evacuation**, **Cleanup evacuation rehearsal**, or the **Regenerate evacuation plan** operation, you can view the workflow details in the **Activities** view.

See "Generating an evacuation plan" on page 151.

See "Regenerating an evacuation plan" on page 152.

See “Performing evacuation” on page 153.

See “Performing rehearse evacuation” on page 153.

See “Performing cleanup evacuation rehearsal” on page 153.

Generating an evacuation plan

Using the Resiliency Platform console you can generate an evacuation plan that lets you evacuate all the assets from the production data center to the recovery data center.

Using the evacuation plan template you can define the sequence in which the virtual business services (VBS) should be migrated from the production data center to the recovery data center. Resiliency groups that do not belong to any VBSs, are appended at the end of the evacuation plan workflow after the VBS.

By default only one priority group is created. To add more priority groups, click **Change Priority** and click the **+** button. You can drag and drop the VBSs into different priority groups.

Reset to Default removes all priority groups except one. All VBSs are moved into a single priority group.

If an asset within a VBS or a resiliency group fails to recover, the evacuation plan skips the asset and continues the process for the remaining assets. To do this you need to select the **Ignore failures** check box while creating the evacuation plan.

If any VBSs and resiliency groups do not fit the evacuation plan criteria, a message is displayed. We recommend that you fix the issues before creating the plan.

Only users with **Manage Evacuation Plans** permission can create and run the evacuation plans.

To generate an evacuation plan

1. Prerequisites

   See “About evacuation plan” on page 149.

2. Navigate

   Automation Plans (menu bar) > **Evacuation Plans**

3. Select **Evacuation Plans**.
4. For the required data center click **Generate Plan**.
5. Review the message if any and click **Next**.
6. Click **Change Priority** if you want to add more priority groups. Click **Submit**.

See “Performing evacuation” on page 153.

See “Performing rehearse evacuation” on page 153.

See “Performing cleanup evacuation rehearsal” on page 153.

See “Regenerating an evacuation plan” on page 152.

Regenerating an evacuation plan

After successfully creating an evacuation plan, if any of the following scenarios occur, you need to regenerate the evacuation plan.

- VBSs are added, modified, or deleted.
- Resiliency groups are added and configured for disaster recovery.
- Existing resiliency group is configured for disaster recovery.

No action is required in the following scenarios:

- Resiliency groups are added.
- Resiliency groups are modified.
- Resiliency groups which are not configured for disaster recovery are deleted.

To add more priority groups to the plan, click **Change Priority** and click the **+** button. You can drag and drop the VBSs into different priority groups. **Reset to Default** removes all priority groups except one. All VBSs are moved into a single priority group.

To regenerate an evacuation plan

1. Navigate

   Automation Plans (menu bar) > **Evacuation Plans**

2. For the required data center click **Regenerate Plan**.
3. Review the message if any and click **Next**.
4. Click **Change Priority** if you want to add more priority groups or click **Reset to Default** if you want to have only one priority group. Click **Submit**.

See “Generating an evacuation plan” on page 151.

See “Performing evacuation” on page 153.

See “Performing rehearse evacuation” on page 153.

See “Performing cleanup evacuation rehearsal” on page 153.

Performing evacuation

Using the Resiliency Platform console, you can run an evacuation plan for a data center which lets you evacuate all the assets from the production data center to the recovery data center.

To run an evacuation plan

1. Navigate

   Automation Plans (menu bar) > **Evacuation Plans**

2. For the required data center, click on the vertical ellipses and select **Evacuate** to run the evacuation plan.

See “Performing rehearse evacuation” on page 153.

See “Performing cleanup evacuation rehearsal” on page 153.

See “Regenerating an evacuation plan” on page 152.

Performing rehearse evacuation

Using the Resiliency Platform console, you can perform a rehearsal of an evacuation plan for a data center. This verifies whether all your assets from the production data center can evacuate to the recovery data center.

To perform a rehearsal of an evacuation plan

1. Navigate

   Automation Plans (menu bar) > **Evacuation Plans**

2. For the required data center, click on the vertical ellipses and select **Rehearse Evacuation**.

See “Performing cleanup evacuation rehearsal” on page 153.

See “Regenerating an evacuation plan” on page 152.

Performing cleanup evacuation rehearsal

After you have performed the rehearse evacuation operation successfully to verify if all your assets from the production data center can evacuate to the recovery data center, you can use the cleanup evacuation rehearsal operation to clean up the rehearsal virtual machines and its volumes in the VBS or resiliency groups.

All temporary objects that are created during the rehearse evacuation operation are now deleted.

During the rehearse evacuation operation, if any virtual machines are in ERROR state, then during the cleanup evacuation rehearsal operation, these virtual machines and their volumes are not deleted. You need to manually delete them. Similarly if the recovery data center is Cloud, then manually delete the instances which are in ERROR state.

To perform the cleanup rehearsal of an evacuation plan

1. Navigate

   Automation Plans (menu bar) > **Evacuation Plans**

2. For the required data center, click on the vertical ellipses and select **Cleanup Evacuation Rehearsal**.

See “Performing evacuation” on page 153.

See “Performing rehearse evacuation” on page 153.

Appendix A

General troubleshooting

This appendix includes the following topics:

- Viewing events and logs in the console
- Events in VMware virtual machines disaster discovery
- Troubleshooting discovery of assets
- Log files to troubleshoot Veritas Resiliency Platform Data Mover
- Managing tunable parameters
- Resiliency Platform fails to attach storage policy to virtual machines
- Resiliency Platform fails to create storage policy
- Resolving the Admin Wait state
- Troubleshooting NetBackup issues
- Troubleshooting delete resiliency group operation

Viewing events and logs in the console

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

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

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

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

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

To view events and logs

1 Navigate

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

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

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

Events in VMware virtual machines disaster discovery

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

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

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

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

Troubleshooting discovery of assets

When asset infrastructure is added to the Infrastructure Management Server (IMS), or when changes are made to the infrastructure, the IMS discovers and correlates the asset information and displays the information on the Assets page of the Resiliency Platform console. The discovery can take some time before the information is updated on the console. Until discovery is complete, some information needed to configure resiliency groups may be missing from the Assets page on the console.

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

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

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

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

Situation	Troubleshooting/best practices
Adding storage arrays as enclosures to IMS	Ensure that the storage arrays that are added to the IMS are the ones that provide storage to the ESX servers managed by the vCenter Server that is added to the IMS.
More than one IMS in data center	Ensure that the vCenter Server that is managing the ESX servers, and the enclosure providing storage to those servers, are added to the same IMS.

Table A-1 Configuring asset infrastructure in IMS for storage arrays in VMware environment *(continued)*

Situation	Troubleshooting/best practices
Refreshing the IMS after a change in infrastructure	Ensure that you use the Refresh operation on the correct vCenter Servers and enclosures where the change was made.
Refreshing the IMS after a change in infrastructure, where there is more than one IMS	Ensure that you use the Refresh operation in the correct IMS.

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

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

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

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

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

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

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

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

Log files to troubleshoot Veritas Resiliency Platform Data Mover

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

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

Issue type	Log files	Location
VMware related operations, specifically, issues related to IO filter installation	/var/opt/VRTSsfmh/logs/vsphere_wsclient_op.log	IMS or Control Host
VMware storage policy related issues	/var/opt/VRTSsfmh/logs/vsphere_spbmclient_pm.log /var/opt/VRTSsfmh/logs/vsphere_spbmclient.log	IMS or Control Host

Table A-4 Log files locations to troubleshoot Resiliency Platform Data Mover issues *(continued)*

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

Managing tunable parameters

Resiliency Platform lets you control some subtasks of an operation by changing certain parameters. Each parameter has a default value which can be overridden. The following table lists the parameters and the conditions in which they can be used.

Table A-5

Parameter name	Description
vmware.hoststorage.unmount.vmfs.timeout	While migrating a resiliency group from one data center to another, the operation may fail to unmount the VMFS datastores on the ESX servers because the resource is still busy. You can increase the wait time for the unmount operation by modifying this parameter. Unit: seconds Default value: 20 Services to restart: wf (workflow)
vmware.hoststorage.unmount.vmfs.retry_interval	While migrating a resiliency group from one data center to another, the operation may fail to unmount the VMFS datastores on the ESX servers because the resource is still busy. The Resiliency Platform retries the unmount operation for a few times before failing the operation. You can change the retry interval by modifying this parameter. Unit: seconds Default value: 5 Services to restart: wf (workflow)

You can modify the default values using the Command Line Interface Shell (klish) menu.

Managing the tunable parameters

1. Logon to klish console as Admin user and then switch to Support user. For support user credentials, you need to contact Veritas support.
2. Start the shell using “shell” command of the klish menu.
3. Open the file `/var/opt/VRTSitrp/conf/itrp.conf` in an editor and modify the required parameters.

   Save the file.

4 Restart the services using the following command.

```
/opt/VRTSitrp/bin/itrpadm service --restart <service name>[,<service name
```

5 Start or retry the failed operation after the service is restarted.

Resiliency Platform fails to attach storage policy to virtual machines

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

Error message: ESX host does not support the virtual hardware configuration of virtual machine. Check the `IMS/CH: /var/opt/VRTSsfmh/logs/vsphere_wsclient_op.log` log file to verify and resolve the issue.

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

Resiliency Platform fails to create storage policy

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

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

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

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

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

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

Resolving the Admin Wait state

During replication of a resiliency group, an issue may occur which requires manual intervention for a Veritas Replication Set on the Replication Gateway. For example, there may be disk corruption or a disk may not be accessible. Any of the replication services, including the I/O Receiver, the Transceiver, the Applier, or the Replication

Engine, may detect an issue when sending, receiving, or applying an update. The operation fails, and the replication state of the Resiliency Group shows in the web console as Admin Intervention Required.

To debug further, determine the type of issue and which Veritas Replication Set is affected by logging into the Replication Gateway and performing step 1 in the procedure described below. A Veritas Replication Set consists of all the disks for a protected virtual machine.

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

To clear the Admin Wait state

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

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

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

   See "Admin Wait state codes " on page 163.

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

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

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

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

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

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

Admin Wait state codes

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

Table A-6 Admin Wait state codes

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

Table A-6 Admin Wait state codes *(continued)*

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

Table A-6 Admin Wait state codes *(continued)*

Admin Wait string	Cause	Resolution	Needs full sync?
Applier: Target Disk not found Error	The target disk is not found on the gateway	If the issue occurred during replication, check if the disk is attached to the cloud Replication Gateway. Also check OpenStack. If the issue occurred during Resycn operation, check that the disk is attached to the on-premises Replication Gateway. Also check iSCSI connection. Check `/var/log/messages` for any disk or virtio or virtblk or iSCSI errors. Check `vxgwaplrd` logs. If the disk was not attached, attach it and clear the admin wait flag. Full synchronization is not required if the update sets were not deleted. Otherwise, full synchronization is required.	Maybe.
Applier: Data Corruption Error	Replication data files are corrupted.	If the files are corrupted, the resolution requires `fsck` and check disk. Perform a full synchronization.	Yes.
IOReceiver: Disk Error	Opening or reading or writing the update set files has failed. Or the replication data directory is full. `/var/opt/VRTSitrpgw/repldata`	Check `/var/log/messages` for any disk errors or file system errors. Check `/var/opt/VRTSitrpgw/log/vxgwiorecvd.log` log for `admin_wait_event` with a `message` value. For example, `"admin_wait_event": 2, "message": "Update set file write error."` Check whether the current update set is accessible. If the replication data directory is full, add a disk using the CLISH menu. For more information on CLISH menu, see the *Deployment guide*. If the files cannot be restored even after manual intervention like mount/fsck/check disk, the disk must be replaced in the configuration and full synchronization is required.	Maybe.

Table A-6 Admin Wait state codes *(continued)*

Admin Wait string	Cause	Resolution	Needs full sync?
IOReceiver: Data Corruption Error	Replication data files are corrupted.	If the files are corrupted, the resolution requires `fsck` and check disk. Perform a full synchronization.	Yes.

Troubleshooting NetBackup issues

The following table lists some of the issues that you may come across while restoring backup images using NetBackup.

Table A-7

Description	Solution
Restore operation fails at the Restore virtual machine image step.	In the **Recent Activities** panel, check the details of the failed step. Details such as Job ID and NetBackup master server name are displayed. These can be used to retrieve the failure reason using `bperror` command on the NetBackup master server.
Although the Restore operation is complete, the Recent Activities panel shows that the process is stuck at "Restore virtual machine image" step.	Do the following: ■ Go to **Settings** > **Infrastructure**, in the details view of the recovery data center > **Copy Manager**. Check the connectivity of NetBackup master server. If it is disconnected, refresh the NetBackup master server. ■ Ensure that there are no time sync issues between Resiliency Platform and NetBackup master server. ■ Check the restore status in the NetBackup master server using the Job ID provided in the details view in **Recent Activities** panel

Table A-7 *(continued)*

Description	Solution
After completion of the Restore operation, the NICs attached to the virtual machines are not displayed.	Ensure that VLAN mapping is complete at the production data center. If multiple network interface cards (NICs) are assigned to a virtual machine, then you need to apply IP customization to all the NICs.
Restore operation fails	Check the log file for errors. Log file location: /var/opt/VRTSsfmh/logs/mhrun.log
Restore operation fails. The activities monitor on NetBackup master server shows the following error. *NetBackup VMware policy restore error (2820)*	When you add a VMware access host or a recovery host to the NetBackup master server, ensure that you add only the NetBackup client. Else the restore operation performed using the Resiliency Platform console may fail.
After adding the NetBackup master server, if you immediately perform the 'Configuring the assets for remote recovery using the Copy service objective' operation, then the operation may fail.	Discovery of NetBackup master server by Resiliency Platform takes some time. If there are no configuration issues, then you may perform the 'Configuring the assets for remote recovery using the Copy service objective' operation after some time.
Deleted NetBackup policies	Resiliency Platform does not validate deactivated NetBackup policies but displays RPO risks if the assets are protected.

Troubleshooting delete resiliency group operation

While performing the delete resiliency group operation, if any of the sub-tasks fail, you can perform the following steps to reclaim the resources.

Perform the following tasks on Replication Gateways on the production and Cloud data center

1. Check for participating consistency groups.

   ```
   /opt/VRTSsfmh/bin/xprtlc -l http://localhost:8080/ConsistencyGroup
   ```

2. Verify that replication is stopped on the gateway.

   ```
   /opt/VRTSsfmh/bin/xprtlc -l
   http://localhost:8080/ConsistencyGroup/<CG_ID>/state
   ```

3. Abort replication on gateway if replication is not stopped.

   ```
   /opt/VRTSsfmh/bin/xprtlc -m POST -l
   http://localhost:8080/ConsistencyGroup/<CG_ID>/abort
   ```

4. Delete the consistency groups.

   ```
   curl -X DELETE
   http://localhost:8080/ConsistencyGroup/<CG_ID>/delete
   ```

Perform the following tasks on the hosts on the production data center

1. Verify that the consistency groups are deleted.

   Linux host:

   ```
   /opt/VRTSitrptap/bin/vxtapinfo status
   ```

   Windows host:

   ```
   C:\Program Files\Veritas\VRTSitrptap\cli\vxtapinfo status
   ```

2. Delete the consistency groups if the state is active.

   - Linux host:

     ```
     /opt/VRTSitrptap/bin/vxtapaction stop -cg <CG_ID>
     /opt/VRTSitrptap/bin/vxtapconfigure delcg -cg <CG_ID> -force
     ```

   - Windows host:

     ```
     C:\Program Files\Veritas\VRTSitrptap
     \cli\vxtapaction stop -cg <CG_ID>
     C:\Program Files\Veritas\VRTSitrptap
     \cli\vxtapconfigure delcg -cg <CG_ID> -force
     ```

3 Verify that the journal disk is removed from each of the virtual machines on the production data center.

Size of the journal disk is usually 1 GB and naming format is:

[<datastore-name>] <vm-hostname>/drl-<uuid>/ITRPSRDRLDisk.vmdk

Remove the journal disk using vSphere client or Hyper-V UI.

Ensure that the file is deleted.

4 Unconfigure network setting on the virtual machines on the production data center.

- Linux host:

```
/opt/VRTSsfmh/bin/perl
 /opt/VRTSsfmh/util/reconfig_vm_settings_on_prim  -unconfigure
```

- Windows host:

```
C:\Program Files\Veritas\VRTSsfmh\bin\perl.exe
 C:\Program Files\Veritas\VRTSsfmh\util
 \reconfig_vm_settings_on_prim -unconfigure
```

5 For physical machines, you need to unsign the DRL disk:

- Linux host:
  Get the device name of DRL disk by executing the following command:

```
/opt/VRTSitrptap/bin/vxtapdrlfind
```

Execute the following command to unsign DRL disk:

```
/opt/VRTSitrptap/bin/vxtapdrlsign clear -drl_dev device_name
```

Where, *device_name* is the device name of the DRL disk.

- Windows host:
  Get the device name of DRL disk by executing the following command:

```
C:\Program Files\Veritas\VRTSitrptap\cli\vxtapdrlfind.exe
```

Execute the following command to unsign DRL disk:

```
C:\Program Files\Veritas\VRTSitrptap\cli\vxtapdrlsign clear
 -drl_dev device_name
```

Where, *device_name* is the device name of the DRL disk.

6 Refresh the virtualization server if the resiliency group contains virtual machines.

Perform the following on cloud virtual machines:

- Verify that the migrated cloud virtual machines are terminated and their volumes are deleted, including the Replication block tracking disk.
  Note that the naming convention for Replication block tracking disk is `DRLVolume_RaaS_<GUID>` whereas for other disks the name starts with `resiliency-group-name_virtual-machine-name`.
- Delete any snapshots for cloud volumes.
- Detach the cloud volumes from the cloud gateway.
- Remove the cloud virtual machines from the cloud IMS.

See "Deleting a resiliency group" on page 81.

Appendix B

Sample policy and trust relationships for AWS

This appendix includes the following topics:

- Sample policy statement for AWS
- Sample trust relationship for AWS

Sample policy statement for AWS

Following is a sample policy statement that you can use to manage the permissions for the users. This policy assigns the following permissions:

- s3:GetBucketLocation and s3:GetObject on vrp-bucket
- ec2:ImportSnapshot, ec2:DescribeSnapshot, and ec2:CopySnapshot on all the resources

```
{
    "Version": "2012-10-17",
    "Statement": [
        {
            "Effect": "Allow",
            "Action": [
                "s3:GetBucketLocation",
                "s3:GetObject"
            ],
            "Resource": [
                "arn:aws:s3:::drlbucketqcowsk",
                "arn:aws:s3:::drlbucketqcowsk/*"
            ]
        },
        {
            "Effect": "Allow",
            "Action": [
                "ec2:ImportSnapshot",
                "ec2:DescribeSnapshots",
                "ec2:CopySnapshot"
            ],
            "Resource": "*"
        }
    ]
}
```

Sample trust relationship for AWS

Following is the sample trust relationship for making the service `vmie.amazonaws.com` assume the role that is associated with the policy.

```
{
  "Version": "2012-10-17",
  "Statement": [
    {
      "Effect": "Allow",
      "Principal": {
        "Service": "vmie.amazonaws.com"
      },
      "Action": "sts:AssumeRole"
    }
  ]
```

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.

klish Command Line Interface SHell. Provides the command line menu on the virtual appliance for use after the initial bootstrap configuration.

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

For the disaster recovery use case, the resiliency domain must contain at least two data centers in different locations, a production data center and recovery data center. Each data center has a Resiliency Manager and one or more IMSs.

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

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

Infrastructure Management Server (IMS) The Veritas Resiliency Platform component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager.

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

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

Index

A
access profile
  configuring for a data center 60
activities
  abort 147
  view 146
asset infrastructure
  troubleshooting discovery of assets 157
asset types
  about 134
assets
  configuring for monitoring 75
  display overview 134
AWS
  support 32

B
Bind settings for data center 62

C
clearing outage 118
configuring for remote recovery
  AWS customization 97
  target asset selection 91
  using 3rd party replication 89
  using AWS 96
  using Data Mover 93
  using vCloud 98

D
dashboard 132
Data Mover
  about 19
  architecture 20
  configuration options 94
  handling application writes 23
  key configuration steps 57
  log files 159
  overview 18
  supported environments 19
Data Mover *(continued)*
  viewing replication states 101
data replication
  handling migration 111
  handling rehearsal 107
  synchronizing data 22
delete resiliency group
  troubleshooting 168
disaster recovery
  configuration 42
  limitations 89
  overview 84
  prerequisites for virtual machines 85
  Resiliency Platform Data Mover 57
  using Resiliency Platform 12
disaster recovery operations
  cleanup rehearsal 110
  migrate 113
  rehearse 106, 108
  rehearse operations 105
  rehearse using NetBackup 109
  resync 115
  takeover 114
DNS server settings for data center 62

E
encryption for data replication 67
evacuation plan
  about 149
  cleanup evacuation rehearsal 153
  evacuating 153
  generating 151
  regenerating 152
  rehearse evacuation 153
events 155–156

F
full synchronization
  about 22

H

HPE 3PAR Remote Copy
- prerequisites 51

I

Infrastructure Management Server
- troubleshooting discovery of assets 157

L

logs
- viewing in console 155

M

migration
- about 111

monitoring
- events 156

N

NetBackup
- about 29
- overview 31
- protecting assets 30

NetBackup Appliances
- about 29

network customization
- options 92
- prerequisites 93

network settings
- configuring for a data center 60
- editing for a data center 61

P

permissions
- about 17

policy statement
- AWS 172

R

Recovery to AWS
- overview 33

Recovery to premises
- Data Mover 24
- third-party replication 27

Recovery to vCloud
- overview 37

rehearse operations 105, 107–110

remote recovery
- using NetBackup images 99

replication
- EMC RecoverPoint 46
- EMC SRDF 44
- Hitachi True copy 50
- IBM SVC Global Mirror 53
- IBM XIV 53
- monitoring events 156
- NetApp SnapMirror 47

Replication Block Tracking disk
- about 23

Replication Gateway
- about 66
- creating pair 67
- modifying encryption scheme 68
- removing pair 69
- viewing 68
- viewing pair 69

replication lag 80

replication status 80

reports
- current risk 143
- historical risk 144
- viewing 145

resiliency groups
- about 72
- deleting 81
- displaying information and status 77
- editing 80
- guidelines for organizing 73
- roles 84
- starting 76
- stopping 77
- viewing detailed information 80

resiliency plan templates
- create 121
- deleting 125
- editing 125
- viewing 126

resiliency plans
- about 120
- create schedule 130
- creating 126
- custom script 123
- delete schedule 130
- deleting 128
- edit schedule 130

resiliency plans *(continued)*
editing 128
executing 128
manual task 122
view schedule 131
viewing 129
Resiliency Platform
capabilities 15
features and components 13
restoring data
using NetBackup 116
risk insight
about 135
risks
current risk report 143
description 137
historical risk report 144
view information 136

S

service objectives
about 73
subnet mapping 64
supported environments
VMware VAIO 19
synchronizing data
about 22

T

trust relationship
AWS 173
tunable parameters
managing 160

V

vCloud
support 36
Veritas Resiliency Platform
about 11
virtual machines
configuring for monitoring 75
VMware using Resiliency Platform
managing 16
VMware vCenter Server privileges 55
vtstap
about 23