Veritas™ Resiliency Platform 3.2 Overview and Planning Guide

VERITAS™

Last updated: 2018-05-30

Document version: Document version: 3.2 Rev 0

Legal Notice

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

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

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

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

doc.feedback@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1

Overview of Resiliency Platform

This chapter includes the following topics:

About Veritas Resiliency Platform

Resiliency Platform has the following core capabilities:

See "About Resiliency Platform features and components" on page 5.

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. See "About Resiliency Manager" on page 7.

Infrastructure Management Server (IMS)	The component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager. The IMS is deployed as a virtual appliance. To achieve scale, multiple IMSs can be deployed in the same data center. See "About Infrastructure Management Server (IMS)" on page 7.
resiliency domain	The logical scope of a Resiliency Platform deployment. It can extend across multiple data centers. See "About resiliency domain" on page 9.
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. If you are using Resiliency Platform Data Mover for replication, each data center must also have at least one Replication Gateway.
asset infrastructure	The data center assets that you add to Resiliency Platform for discovery and monitoring by the IMS.
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 perform the disaster recovery operations on the entire VBS.

About Resiliency Manager

The Resiliency Manager includes a set of loosely coupled services, a distributed data repository, and a management web console. The Resiliency Manager provides the services required for protecting assets, such as virtual machines, within a resiliency domain. It also provides workload automation services.

The Resiliency Manager discovers and manages information about data center assets from an Infrastructure Management Server (IMS), which is another required Resiliency Platform component. The Resiliency Manager stores the asset information in its data repository and displays the information in its management console.

See “About resiliency domain” on page 9.

See “About Infrastructure Management Server (IMS)” on page 7.

About Infrastructure Management Server (IMS)

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

The IMS sends information about the assets to the Resiliency Manager so that the Resiliency Manager can manage the assets. Management operations on assets (for example, starting or stopping virtual machines) that you initiate from the web console are carried out by the IMS.

You can also configure multiple Infrastructure Management Servers in the same data center. For example, to achieve scale, you can add a separate IMS for a separate business unit such as Human Resources or Finance. More than one IMS can be managed by the same Resiliency Manager.

See “About resiliency domain” on page 9.

See “About Resiliency Manager” on page 7.

About Replication Gateways

If you plan to use Resiliency Platform Data Mover for replication of data in your environment, you need to deploy and configure at least one Replication Gateway in your production as well as recovery data center.

The Replication Gateway component of Veritas Resiliency Platform is a staging server that aggregates and batches data from multiple virtual machines during replication. The Gateway also performs data optimization like local deduplication and compression. The Gateway on production data center is always paired with a Gateway on recovery data center. The recovery data center Gateway is a staging server that applies the data on the recovery data center storage.

Each Replication Gateway includes the following components:

- I/O receiver
  Receives the application I/Os that were tapped and sent by the application host 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 cloud Gateway.
- Scheduler
  Manages the jobs and policies in the Gateway.
- Engine
  Maintains the state of replication and also coordinates with all other components.

About Data Gateway

If you want to choose Object Storage replication mode for migration of your assets to AWS, you need to deploy a Data Gateway in AWS environment.

The Data Gateway acts like a communication channel between the on-premises Replication Gateway and cloud Replication Gateway. The data being replicated from the on-premises data center gets compressed and stored in S3 bucket in the form of objects. The cloud Replication Gateway pulls this data from S3 bucket, decompresses it and applies to the target disk.

You can use a single Data Gateway for replicating data between multiple Replication Gateways.

To deploy the Data Gateway in AWS, you need to download a zip file that is shipped along with Veritas Resiliency Platform.

A few resources get created in AWS when you deploy a Data Gateway in the AWS environment. You must not delete these resources while the Data Gateway is in use as it may impact the functionality of the feature and the product. These resources automatically get deleted when you delete the Data Gateway.

About resiliency domain

The resiliency domain is a logical object that you create from the web console after you deploy the Resiliency Manager.

For disaster recovery, the resiliency domain must contain at least two data centers, a production data center and a recovery data center that can be on-premises or in the cloud.

A resiliency domain can optionally be implemented at a single data center for automation of workload tasks.

See “About Resiliency Manager” on page 7.

See “About Infrastructure Management Server (IMS)” on page 7.

Chapter 2

Planning your Resiliency Platform environment

This chapter includes the following topics:

Replication in a Resiliency Platform deployment

Veritas Resiliency Platform supports several forms of replication for data recovery from your production data center to your recovery data center.

- Array-based replication (block-based replication) using supported arrays
- Hypervisor-based replication using Hyper-V Replica
- NetBackup Auto Image Replication (AIR)
- Resiliency Platform Data Mover (separately licensable feature of Resiliency Platform)

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

Figure 2-1 Replication in a Resiliency Platform deployment

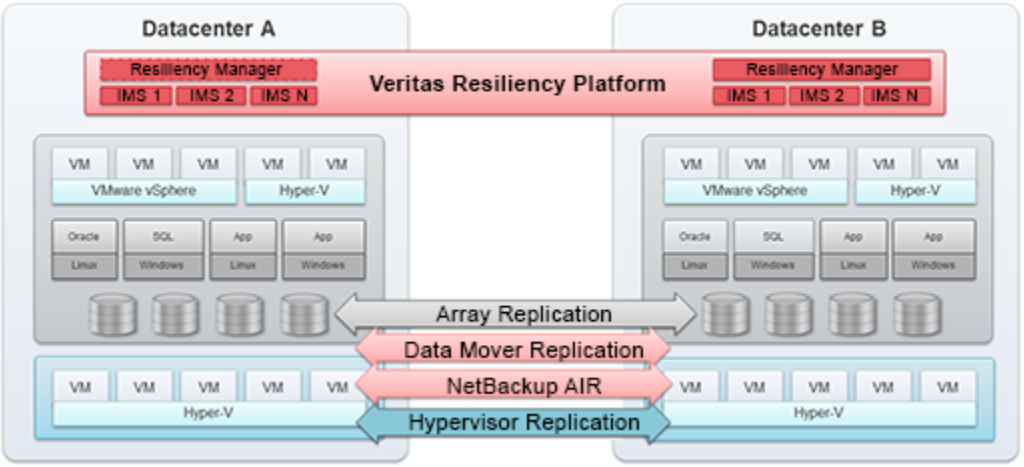

About Veritas Resiliency Platform Data Mover

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.

How Resiliency Platform Data Mover works

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.

About replication tunables

If you are using Resiliency Platform Data Mover to replicate your data across data centers, you have the option to tune some replication parameters to optimize the performance and scalability of the replication.

Resiliency Platform Data Mover replicates data from the protected virtual machines at the source data center to the target disks at target data center. The Replication Gateway component of Veritas Resiliency Platform is a staging server that aggregates and batches data from multiple virtual machines during replication. The staging storage on the Replication Gateway component of Veritas Resiliency Platform consists of two sections:

- **Reserve storage**: Replication Gateway allocates a certain part of the staging storage disk to each of the protected virtual machine. This part of the staging storage is called reserve storage.
- **Shared pool**: This part of the staging storage is shared among all the virtual machines.

Parameters that can be tuned

Following are the three parameters for Replication Gateway that you can tune to optimize the Recovery Point Objective (RPO), replication performance, and scalability:

- **Update set**: Before sending date to the target Replication Gateway, the source Replication Gateway implements a data optimization technique and creates a set of data. This set of data collected over a period is called an update set.
- **Replication Frequency**: The period set to delimit and cut the update set is called replication frequency. Once an update set is cut, it gets scheduled for transfer to the target Replication Gateway.
- **Quota per Veritas Replication Set**: The space reserved for each virtual machine in the reserve storage is called quota per Veritas Replication Set.

Table 2-1

Tunable type	Change impact
Quota per Veritas Replication Set (Quota-per-CG)	Scale (number of virtual machines protected by the Replication Gateway)
size of update set (Update-set-size)	Performance (local deduplication), compression, RPO)
Replication frequency	RPO for all the resiliency groups configured on the gateway

The following formulae are used to link above replication tunables:

Reserve storage = number of virtual machines * Quota per Veritas Replication Set

Number of update sets = Quota per Veritas Replication Set / size of update set. The space for these many update sets is always reserved on the gateway.

You can tune the Quota-per-CG, update-set-size, and replication-frequency using the klish menu commands.

About Veritas Resiliency Platform Data Mover architecture

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.

About synchronization using Veritas Resiliency Platform Data Mover

Veritas Resiliency Platform Data Mover uses two types of synchronization techniques for replicating the data from source to target data center:

- See "About full synchronization of data" on page 13.
- See "About incremental synchronization" on page 14.

About full synchronization of data

Veritas Resiliency Platform uses full synchronization only in the following conditions:

- After disaster recovery configuration for a resiliency group:
  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 a full synchronization.
- During Resync operation performed after a takeover operation:
  A full synchronization is also required after a takeover. Takeover is an activity initiated by a user when the source data center is down due to a disaster, and

the virtual machines need to be brought up at the target (recovery) data center to provide business continuity. 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. This Resync operation launches a full synchronization to synchronize the data on the production data center with the data in the recovery data center. When the synchronization completes, the production data center is up-to-date. You can then perform the Migrate operation.

- After addition or removal of a disk from any of the protected virtual machines: If you add or remove a disk from any of the protected virtual machines, a risk is raised. You need to edit the resiliency group to resolve this risk. During this edit resiliency group operation, you remove the affected virtual machine. Edit the resiliency group one more time to add the virtual machine again and update the configuration. Full synchronization is launched after your edit resiliency group operation gets completed.

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 bandwidth of the LAN and WAN environment, and the amount of I/O 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.

If the replication state is Syncing, you can view the status of data replication on the resiliency group details page. The progress is displayed on a status bar with percentage complete information. Time required to sync the data is also displayed on this page.

At times, you need to manually invoke a full synchronization to resume replication after a disk failure or infrastructure failure. For more information on conditions where a full synchronization is required:

See “About synchronization using Veritas Resiliency Platform Data Mover” on page 13.

About incremental synchronization

An incremental synchronization targets to synchronize only that data which has changed since the last synchronization (either incremental or full synchronization). Incremental synchronization saves much of the time and resources used in replication of the data between the data centers.

Except the two conditions where a full synchronization is performed in veritas Resiliency Platform (after disaster recovery configuration and after a takeover operation), at all other times, Resiliency Platform uses incremental synchronization while replicating the data from source data center to target data center. These

instances where incremental synchronization is used in Resiliency Platform include the following:

- If there are any network failures in the replication path
- If there is a system reboot of Replication Gateway or protected virtual machines
- If you replace a healthy or faulted Replication Gateway with another Replication Gateway
- If you perform a migrate operation. In this case, the virtual machines are brought up on the target site and then direction of replication changes. At this point, Resiliency Platform uses incremental synchronization.

If the replication state is Syncing, you can view the status of data replication on the resiliency group details page. The progress is displayed on a status bar with percentage complete information. Time required to sync the data is also displayed on this page.

See "About synchronization using Veritas Resiliency Platform Data Mover" on page 13.

How Veritas Resiliency Platform Data Mover handles virtual machine writes

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.
- IO is written to virtual machine storage.
- 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.
- IO is written to virtual machine storage.
- 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.

About Object Storage mode replication

If you want to migrate your data center assets to AWS, you have an option to choose between Object Storage mode replication and Direct mode replication.

In Resiliency Platform, the Object Storage mode replication is used to leverage the S3 Object Storage services provided by AWS. The following are some of the advantages of using the Object Storage mode replication in Resiliency Platform:

- Automatically scales according to the requirements of the user by utilizing the AWS services to achieve scalability.
- Facilitates resiliency for the Replication Gateway. Since the data keeps getting replicated and stored in S3 bucket, failure of replication gateway in the cloud does not hamper the replication.

To enable the Object Storage mode replication, you need to deploy a Data Gateway in AWS environment.

See "About Data Gateway" on page 8.

Recovery options using Resiliency Platform

There are various recovery options available with Resiliency Platform. You can use any of the supported third-party replication technologies, NetBackup, or Resiliency Platform Data Mover to replicate and recover your data across data centers. You can also recover your InfoScale applications using Resiliency Platform.

Table 2-2

Category	Recovery options
Using third-party replication	Recovery to on-premises data center: ■ Recovery of VMware virtual machines to on-premises data center ■ Recovery of Hyper-V virtual machines to on-premises data center ■ Recovery of applications to on-premises data center

Table 2-2 *(continued)*

Category	Recovery options
Using NetBackup	Recovery to local and remote data center: ■ Recovery of VMware virtual machines to local and remote data center
Using Resiliency Platform Data Mover	■ Recovery to on-premises data center: ■■ Recovery of VMware virtual machines to on-premises data center using VAIO ■ Recovery to AWS data center: ■■ Recovery of VMware virtual machines to AWS data center ■■ Recovery of Hyper-V virtual machines to AWS data center ■ Recovery to vCloud data center: ■■ Recovery of VMware virtual machines to vCloud data center ■■ Recovery of Hyper-V virtual machines to vCloud data center ■■ Recovery of VMware virtual machines to vCloud data center without adding the vCenter Server ■■ Recovery of Hyper-V virtual machines to vCloud data center without adding the Hyper-V Server ■■ Recovery of virtual machines from vCloud Director to vCloud Director ■ Recovery to Azure data center: ■■ Recovery of VMware virtual machines to Azure data center ■■ Recovery of Hyper-V virtual machines to Azure data center ■ Recovery to OpenStack data center: ■■ Recovery of VMware virtual machines to OpenStack data center ■■ Recovery of Hyper-V virtual machines to OpenStack data center
Using Veritas InfoScale Management Server	Recovery to on-premises data center: ■ Recovery of InfoScale applications to on-premises data center

Planning a resiliency domain for efficiency and fault tolerance

Before you deploy Veritas Resiliency Platform, you should plan how to scale the deployment for efficiency and fault tolerance.

The source and target data centers do not require a one-on-one mapping of IMSs. For example, you can have two IMSs in the source data center and one IMS in the target data center.

Veritas Resiliency Platform enables you to provision for resiliency of its components. Following are some ways to achieve the resiliency of Resiliency Platform components.

Table 2-3 Recovery of Resiliency Platform components from disaster scenarios

Resiliency Platform component	Steps to recover
Resiliency Manager	See "Recovering Resiliency Manager" on page 18.
Infrastructure Management Server (IMS)	See "Recovering IMS" on page 19.
Replication Gateway	See "Recovering Replication Gateway" on page 21.

See "About resiliency domain" on page 9.

See "About Resiliency Manager" on page 7.

See "About Infrastructure Management Server (IMS)" on page 7.

Simplified trialware deployment experience

Veritas Resiliency Platform 3.2 introduces a simplified trialware or Proof of Concept (POC) deployment experience for a new user. You get use case based resources that include zip files to be used for deployment and configuration, a configuration planner, and a Quick Start guide explaining the entire flow along with the documentation link.

Recovering Resiliency Manager

Following are the conditions in which Resiliency Manager may go offline, along with its impact, and steps to recover from such scenario. The table also describes the mitigation steps for each of the condition mentioned.

Table 2-4 Recovering Resiliency Manager

Scenario	Impact	Steps to recover	Mitigation steps
Single Resiliency Manager deployed in cloud and Resiliency Manager goes offline due to public cloud outage.	▪ Resiliency domain goes offline. ▪ You cannot perform any disaster recovery operation from source site to target site or vice versa.	Wait till the cloud provider fixes the cloud outage. There is no action required from Resiliency Platform side. Business is resumed as usual after Resiliency Manager is brought online.	You can deploy multiple (a minimum of three) Resiliency Managers in a data center to achieve resiliency of Resiliency Manager. If your cloud data center is in AWS, It is recommended that you deploy multiple Resiliency Managers in different Availability Zones to achieve maximum resiliency of Resiliency Manager.
One Resiliency Manager deployed per site and any one of the Resiliency Managers goes offline.	Resiliency of Resiliency Managers gets impacted, no impact on product functionality	Bring the affected Resiliency Manager online. Use this recovered Resiliency Manager to control the assets. If Resiliency Manager is in irrecoverable state then perform the Leave domain operation from the other Resiliency Manager in the domain. Deploy a new Resiliency Manager on the impacted site.	Ensure appropriate redundancy of Resiliency Manager against the Hypervisor, storage, and network level failure. You can deploy multiple (a minimum of three) Resiliency Managers in a data center to achieve resiliency of Resiliency Manager.

See “Planning a resiliency domain for efficiency and fault tolerance” on page 17.

Recovering IMS

Following are the conditions in which Infrastructure Management Server (IMS) may go offline, along with its impact, and steps to recover from such scenario. The table also describes the mitigation steps for each of the condition mentioned.

Table 2-5 Recovering IMS

Scenario	Impact	Steps to recover	Mitigation steps
IMS at source data center goes offline	Cannot perform any operation that impacts the source data center configuration. You cannot perform migrate, but can perform rehearsal, rehearsal cleanup, and takeover operations.	■ If IMS is brought online: ■ Refresh the hypervisor discovery ■ Refresh the host discovery for the hosts that are attached to the IMS ■ If IMS is in irrecoverable state: See "Steps to recover an IMS" on page 20.	Ensure appropriate redundancy measures for IMS at hypervisor and storage level
IMS at target data center goes offline	Cannot perform any operation that impacts the target data center configuration. You cannot perform migrate, rehearsal, rehearsal cleanup, and takeover operations.	■ If IMS is brought online: ■ Refresh the hypervisor discovery ■ Refresh the host discovery for the hosts that are attached to the IMS ■ If IMS is in irrecoverable state: See "Steps to recover an IMS" on page 20.	Ensure appropriate redundancy measures for IMS at hypervisor and storage level

See "Planning a resiliency domain for efficiency and fault tolerance" on page 17.

Steps to recover an IMS

Following are the steps that you need to perform if an Infrastructure Management Server (IMS) is in irrecoverable state:

- Deploy and configure a new IMS in the data center.
- Move the Replication Gateway, virtualization server, and virtual machines from the faulted IMS to the new IMS.

- An existing Windows Install host can not be moved to the new IMS. For Windows hosts, add a new Windows Install host to the newly configured IMS.
- Similarly, in case of third-party replication, a new storage discovery host needs to be added to the newly configured IMS.

See "Recovering IMS" on page 19.

Recovering Replication Gateway

Following are the conditions in which Replication Gateway may go offline, along with its impact, and steps to recover from such scenario. The table also describes the mitigation steps for each of the condition mentioned.

Table 2-6 Recovering Replication Gateway

Scenario	Impact	Steps to recover	Mitigation steps
Replication Gateway at source site goes offline or unreachable	Replication from source gateway to target gateway stops. You cannot perform migrate operation, but takeover, rehearsal, and cleanup rehearsal operations can still be performed.	▪ If Network gets recovered or gateway is brought online: Replication will resume automatically ▪ If gateway is irrecoverable: Replace gateway	Ensure appropriate network redundancy
Replication Gateway at target site goes offline or unreachable	Replication from source gateway to target gateway stops. You cannot perform ay disaster recovery operation such as migrate, takeover, rehearsal, and cleanup rehearsal.	▪ If Network gets recovered or gateway is brought online: Replication will resume automatically ▪ If gateway is irrecoverable: Replace gateway	Ensure appropriate network redundancy

Note: If you are using Resiliency Platform Data Mover to recover VMware virtual machines to on-premises data center, then you cannot replace the gateway.

See "Planning a resiliency domain for efficiency and fault tolerance" on page 17.

On-boarding with Resiliency Platform

The following table describes the various steps that are involved in the customer on-boarding with Resiliency Platform and what to expect during each of these steps:

Table 2-7 On-boarding with Resiliency Platform

Step	Description

Deploy	■ Deploy Resiliency Platform virtual appliances and configure them as Resiliency Manager, Infrastructure Management Server (IMS), or Replication Gateway
	■ Define the resiliency domain through Getting Started wizard
	■ Add assets to your resiliency domain for discovery:
	  ■ Virtual machines
	  ■ Applications
	  ■ Storage enclosures

Discover	■ Resiliency Platform's deep discovery enables identification of the following:
	  ■ Virtual machines
	  ■ Applications
	  ■ Storage enclosures
	  ■ Software/hardware replication
	  ■ Virtual networks (vSwitches)

Define service level objective	■ Configure service level objective based on the intended Recovery Point Objective (RPO). service level objective driven configuration enables the following capabilities:
	  ■ Basic monitoring of assets
	  ■ Recovery of assets
	  ■ Recovery of multi-tier business services (VBS)
	  ■ Health status reporting of assets
	  ■ Risk alerts and notifications for protected assets

Table 2-7 On-boarding with Resiliency Platform *(continued)*

Step	Description
Manage	■ Single-click rehearsal for resiliency groups and VBS validates disaster readiness: ■ Automated rehearsal ■ Automated rehearsal cleanup ■ Option of network isolation for workloads during rehearsal ■ Single-click recovery or migration of resiliency groups and VBS: ■ Automated recovery or migration based on the service level objective ■ Recovery with predefined network customization ■ Recovery based on predefined grouping or order ■ Controlled recovery using Resiliency Plans ■ Single-click evacuation plan for resiliency groups and VBS: ■ Option of defining priority levels for VBS ■ Automated rehearsal or cleanup rehearsal for evacuation plan

Index