

# Veritas™ Resiliency Platform 3.1 Overview and Planning Guide

Applicable for Veritas Resiliency  
Platform 3.1 and 3.1 Update 1



# Veritas Resiliency Platform: Overview and Planning Guide

Last updated: 2018-02-09

Document version: Document version: 3.1 Rev 1

## Legal Notice

Copyright © 2018 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

Veritas Technologies LLC

500 E Middlefield Road  
Mountain View, CA 94043

<http://www.veritas.com>

## Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

[CustomerCare@veritas.com](mailto:CustomerCare@veritas.com)

Japan

[CustomerCare\\_Japan@veritas.com](mailto: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](mailto: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](https://sort.veritas.com/data/support/SORT_Data_Sheet.pdf)

# Contents

Chapter 1	Overview of Resiliency Platform .....	5
	About Veritas Resiliency Platform .....	5
	About Resiliency Platform features and components .....	6
	About Resiliency Manager .....	8
	About Infrastructure Management Server (IMS) .....	9
	About Replication Gateways .....	10
	About Data Gateway .....	11
	About resiliency domain .....	12
Chapter 2	Planning your Resiliency Platform environment .....	13
	Replication in a Resiliency Platform deployment .....	13
	About Direct mode replication .....	14
	About Object Storage mode replication .....	14
	Recovery options using Resiliency Platform .....	15
	About synchronization using Veritas Resiliency Platform Data Mover .....	16
	About full synchronization .....	16
	About incremental synchronization .....	17
	About update sets .....	18
	When do you need to change the size of the update set .....	18
	Changing the size of the update set .....	19
	Planning a resiliency domain for efficiency and fault tolerance .....	19
	Resiliency of Resiliency Platform components .....	20
	On-boarding with Resiliency Platform .....	21
Index .....		23

# Overview of Resiliency Platform

This chapter includes the following topics:

- [About Veritas Resiliency Platform](#)
- [About Resiliency Platform features and components](#)

## 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 on-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.
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 6.

## 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	<p>The component that provides resiliency capabilities within a resiliency domain. It is composed of loosely coupled services, a distributed data repository, and a management console. The Resiliency Manager is deployed as a virtual appliance.</p> <p>See <a href="#">“About Resiliency Manager”</a> on page 8.</p>
Infrastructure Management Server (IMS)	<p>The component that discovers, monitors, and manages the asset infrastructure within a data center. The IMS transmits information about the asset infrastructure to the Resiliency Manager. The IMS is deployed as a virtual appliance.</p> <p>To achieve scale, multiple IMSs can be deployed in the same data center.</p> <p>See <a href="#">“About Infrastructure Management Server (IMS)”</a> on page 9.</p>

Veritas InfoScale Operations Manager Management Server	<p>The component that allows discovery of InfoScale applications that are already configured in Veritas InfoScale Operations Manager, also referred to as Veritas InfoScale Operations Manager server. You can manage the InfoScale applications that are already configured in Veritas InfoScale Operations Manager on Linux, Solaris, AIX as well as Windows platform.</p> <p>You need to add this component only if you want to manage and recover the Infoscale applications that are already configured in Veritas InfoScale Operations Manager.</p>
NetBackup Server	<p>The component that allows restoration of virtual machines to a local or remote data center using NetBackup generated backup images.</p> <p>You need to add this component only if you want to restore the virtual machines using NetBackup generated backup images.</p>
Replication Gateway	<p>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.</p> <p>If you plan to use any third party replication technology, you do not need to deploy Replication Gateway.</p> <p>See <a href="#">“About Replication Gateways”</a> on page 10.</p>
Data Gateway	<p>The component of Veritas Resiliency Platform that is deployed in AWS data center to enable replication using Object Storage. You need to deploy this component only if you plan to use Object Storage replication for one-time migration to AWS data center.</p> <p>See <a href="#">“About Data Gateway”</a> on page 11.</p>
resiliency domain	<p>The logical scope of a Resiliency Platform deployment.</p> <p>It can extend across multiple data centers.</p> <p>See <a href="#">“About resiliency domain”</a> on page 12.</p>
data center	<p>For a disaster recovery use case, the resiliency domain must contain at least two data centers in different locations, a production data center and recovery data center. Each data center has a Resiliency Manager and one or more IMSs. If you are using Resiliency Platform Data Mover for replication, each data center must also have at least one Replication Gateway.</p>

asset infrastructure	<p>The data center assets that you add to Resiliency Platform for discovery and monitoring by the IMS.</p> <p>The asset infrastructure can include hosts (Windows or Linux servers), virtualization servers for Hyper-V and VMware, and enclosures (storage arrays). Once the asset infrastructure is discovered by the IMS, the discovered virtual machines or applications are listed in the console as assets to manage or protect.</p>
resiliency group	<p>The unit of management and control in Resiliency Platform. You organize related assets into a resiliency group and manage and monitor them as a single entity.</p>
service objective	<p>A template to define the type of operations and technologies that are supported for a group of assets. You apply a service objective to each resiliency group.</p> <p>A template which identifies the characteristics of a service. These could be availability related characteristics such as local redundancy, and number of nodes in a cluster or DR characteristics such as remote recovery, Recovery Point Objective (RPO) SLAs, rehearsal support etc. Service objective is applied when a group of assets are being added to a resiliency group.</p> <p>Resiliency Platform monitors the resiliency groups based on the service objective definition and raises the risks as applicable.</p>
Virtual Business Service (VBS)	<p>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.</p>

## 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.

You start by deploying one Resiliency Manager and creating the resiliency domain. You can then add more Resiliency Managers to the resiliency domain for efficiency

of local access and for fault tolerance. You can deploy multiple Resiliency Managers in the same data center or in separate geographical locations.

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.

Multiple Resiliency Managers that are part of the same domain synchronize their databases using built-in replication. Each Resiliency Manager has its own web console but because the data is synchronized, all consoles show the same data. Operations can be performed from any console and the results show in all the consoles in the resiliency domain.

See [“About resiliency domain”](#) on page 12.

See [“About Infrastructure Management Server \(IMS\)”](#) on page 9.

## 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 asset infrastructure can include objects such as hosts, virtualization servers, and enclosures (storage arrays).

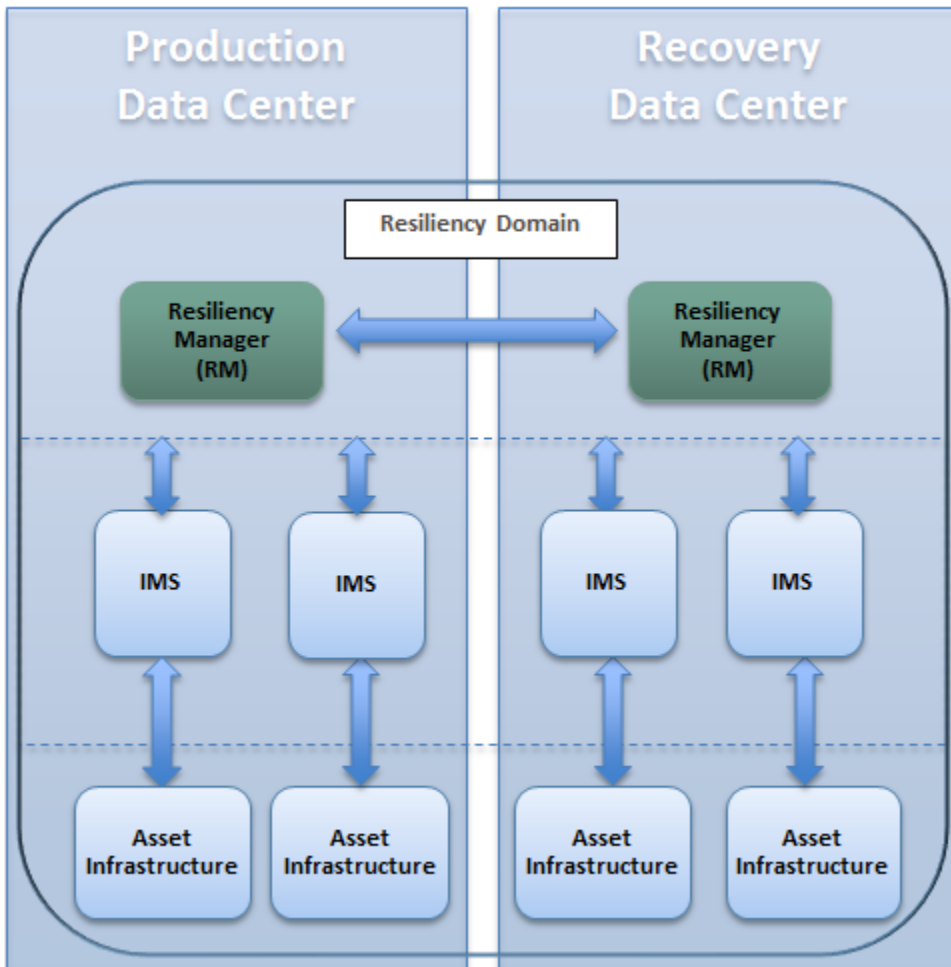
The IMS sends information about the assets to the Resiliency Manager so that the Resiliency Manager can manage the assets. Management operations on assets (for example, starting or stopping virtual machines) that you initiate from the web console are carried out by the IMS.

If there are multiple data centers in different geographical locations, a separate IMS is deployed and configured for each geographical data center location.

Each IMS connects to only one Resiliency Manager at a time. If a Resiliency Manager failure occurs, an IMS can automatically connect to another Resiliency Manager within the same domain.

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.

**Figure 1-1** Multiple Infrastructure Management Servers in a data center



See [“About resiliency domain”](#) on page 12.

See [“About Resiliency Manager”](#) on page 8.

## 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 write cancellation. 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.

See [“About Resiliency Manager”](#) on page 8.

See [“About Infrastructure Management Server \(IMS\)”](#) on page 9.

## 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

A resiliency domain is the management domain of a Veritas Resiliency Platform deployment. It represents the scope of the deployment, which can spread across multiple data centers and can include multiple Resiliency Managers and other components, along with the infrastructure that is being managed and protected. Within the resiliency domain, Resiliency Platform can protect assets, for example, virtual machines and applications, and orchestrate automation of workload tasks for the assets.

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 8.

See [“About Infrastructure Management Server \(IMS\)”](#) on page 9.

# Planning your Resiliency Platform environment

This chapter includes the following topics:

- [Replication in a Resiliency Platform deployment](#)
- [Recovery options using Resiliency Platform](#)
- [About synchronization using Veritas Resiliency Platform Data Mover](#)
- [About update sets](#)
- [Planning a resiliency domain for efficiency and fault tolerance](#)
- [On-boarding with Resiliency Platform](#)

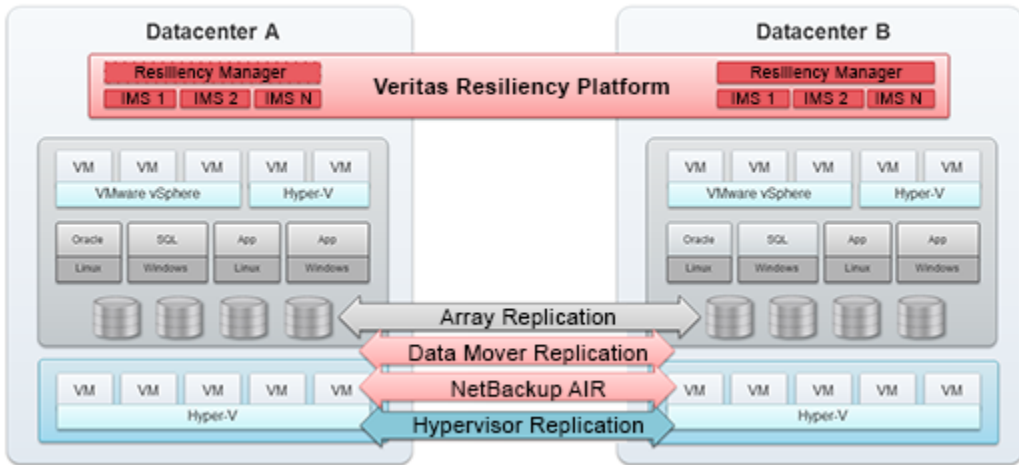
## Replication in a Resiliency Platform deployment

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 Direct mode replication

For migrating your data centre assets to Amazon Web Services (AWS) using Veritas Resiliency Platform, you have an option to choose between two modes of replication: Direct mode and Object Storage mode.

For migrating to a data center other than AWS using Veritas Resiliency Platform, only Direct mode replication is used.

The Replication Gateway on production data center is always paired with a Replication Gateway on recovery data center. In Direct mode replication, the production site Replication Gateway acts as a staging server that aggregates and batches data from multiple virtual machines during replication. The recovery data center Gateway is a staging server that applies the data on the recovery 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 11.

## 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 2-1**

Category	Recovery options
Using third-party replication	Recovery to on-premises data center: <ul style="list-style-type: none"><li>■ Recovery of VMware virtual machines to on-premises data center</li><li>■ Recovery of Hyper-V virtual machines to on-premises data center</li><li>■ Recovery of applications to on-premises data center</li></ul>
Using NetBackup	Recovery to local and remote data center: <ul style="list-style-type: none"><li>■ Recovery of VMware virtual machines to local and remote data center</li></ul>

**Table 2-1** (continued)

Category	Recovery options
Using Resiliency Platform Data Mover	<ul style="list-style-type: none"> <li>■ Recovery to on-premises data center: <ul style="list-style-type: none"> <li>■ Recovery of VMware virtual machines to on-premises data center using VAIO</li> </ul> </li> <li>■ Recovery to AWS data center: <ul style="list-style-type: none"> <li>■ Recovery of VMware virtual machines to AWS data center</li> <li>■ Recovery of Hyper-V virtual machines to AWS data center</li> </ul> </li> <li>■ Recovery to vCloud data center: <ul style="list-style-type: none"> <li>■ Recovery of VMware virtual machines to vCloud data center</li> <li>■ Recovery of Hyper-V virtual machines to vCloud data center</li> <li>■ Recovery of VMware virtual machines to vCloud data center without adding the vCenter Server</li> <li>■ Recovery of Hyper-V virtual machines to vCloud data center without adding the Hyper-V Server</li> </ul> </li> <li>■ Recovery to Azure data center: <ul style="list-style-type: none"> <li>■ Recovery of VMware virtual machines to Azure data center</li> <li>■ Recovery of Hyper-V virtual machines to Azure data center</li> </ul> </li> </ul>
Using Veritas InfoScale Management Server	<p>Recovery to on-premises data center:</p> <ul style="list-style-type: none"> <li>■ Recovery of InfoScale applications to on-premises data center</li> </ul>

## 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”](#) on page 16.
- See [“About incremental synchronization”](#) on page 17.

### About full synchronization

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.

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 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.

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 16.

## 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.

See [“About synchronization using Veritas Resiliency Platform Data Mover”](#) on page 16.

## About update sets

Resiliency Platform Data Mover replicates data from the protected virtual machines at the source data center to the target disks at target data center. To replicate this data, the Data Mover driver taps the workload I/Os and sends it to the source Replication Gateway. The source Replication Gateway implements a data optimization technique and creates a set of workload I/Os. This set of workload I/Os collected over a period of time is called an update set.

The update set is stored on the staging storage disk in its optimized format. While sending this update set over WAN, some other data optimization techniques are used. As a result of this data optimization, the amount of data sent over WAN is much lower than the amount of data processed for replication by the Replication Gateway.

See [“When do you need to change the size of the update set”](#) on page 18.

See [“Changing the size of the update set”](#) on page 19.

## When do you need to change the size of the update set

The default size of update set is tuned to solve most of the applications having moderate sizes of the disks and moderate throughput. These values are tuned to achieve service level agreement (RPO or RTO) for the resiliency group.

You may consider changing the size of the update size in the following situations:

- If you want to protect more virtual machines without allocating more staging storage on the Replication Gateway.
- If the size of the staging storage in the customer environment is smaller than expected.
- If the workload I/O pattern is very low.

Note that changing these values to a value that does not fit your I/O pattern might violate the service level agreement. For example, if the size of the update set is changed to a considerably smaller size and workload I/O rate is very high, the Replication Gateway tries to optimize as much as possible. But if the WAN bandwidth is low, then Replication Gateway may not be able to provide the desired RPO. In

such cases, the replication goes into bit tracking mode. Once enough space is available on the Replication Gateway, replication is automatically resumed.

So, it is expected to have a WAN link with sufficient bandwidth to sustain the workload I/O traffic. Also, to achieve the best optimization possible, keep the update set size to the default value. In cases where workload I/O rate is not high and virtual machine disk size is, you can change the size of the update set and set it to a lower value

If you decrease the size of the update set, the number of virtual machines to be protected increases and vice versa.

See [“About update sets”](#) on page 18.

## Changing the size of the update set

You can change the size of the update set only if there are no resiliency groups configured with this Replication Gateway. After the Replication Gateway is provisioned, you can change the size of the update set by running a klish command on the Replication Gateway.

### To change the size of the update set

- 1 Log into klish as admin user on the Replication Gateway on the source site.
- 2 Run the following command:

```
manage>datamover operation modify-updateset-size
```

The command prompts to enter the size that you want to set for the update set. The default update set size is 2000 MB. The minimum update set size allowed is 500 MB and the maximum update set size allowed is 3000 MB.

---

**Note:** You need to have the same size of update set size on all the peer gateways.

---

See [“About update sets”](#) on page 18.

## 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.

Although a resiliency domain requires only one Resiliency Manager, you can add multiple Resiliency Managers instances to the domain. For example, you can

distribute Resiliency Managers geographically for efficiency of local access. For resiliency, you can even have multiple Resiliency Managers in one data center.

The recommended minimum deployment for disaster recovery to premises data center would be four virtual appliances: a Resiliency Manager and Infrastructure Management Servers (IMS) in the production data center and a Resiliency Manager and IMS in the recovery data center.

The recommended minimum deployment for disaster recovery to cloud data center would be three virtual appliances: an IMS in the production data center and a Resiliency Manager and IMS in the recovery data center.

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

You can add multiple Infrastructure Management Servers (IMS) to a resiliency domain. For example, if there are multiple data centers in different geographical locations to be managed, you configure a separate IMS for each geographical data center location. You can also configure more than one IMS in the same data center.

If you plan to use Resiliency Platform data mover for replication, then additionally you need minimum one Replication Gateway in each data center. Resiliency Platform supports 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.

See [“About resiliency domain”](#) on page 12.

See [“About Resiliency Manager”](#) on page 8.

See [“About Infrastructure Management Server \(IMS\)”](#) on page 9.

## Resiliency of Resiliency Platform components

Following is the description of steps to be performed to achieve the resiliency of Resiliency Platform components.

Table 2-2

Resiliency Platform component	Impact if the component is faulted	Steps to recover
Resiliency Manager	It is recommended to have multiple Resiliency Managers in a data center to achieve resiliency of Resiliency Manager.	If one of the Resiliency Managers gets faulted, you can use the other Resiliency Manager in the same data center.
Infrastructure Management Server (IMS)	You may not be able to perform disaster recovery operations until you replace the faulted IMS with a new IMS.	From Resiliency Platform 3.1 onwards, you have an option to move the Replication Gateway, Virtualization server, and virtual machines from the faulted IMS to the new IMS.
Replication Gateway	Replication stops, you need to replace the gateway to resume the replication.	From Resiliency Platform 3.0 onwards, a single click operation is provided in the console to replace the faulted Replication Gateway using in-guest replication.  For VAIO based replication, you have a manual workaround to replace the Gateway.

## 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 2-3** On-boarding with Resiliency Platform

Step	Description
Deploy	<ul style="list-style-type: none"> <li>■ Deploy Resiliency Platform virtual appliances and configure them as Resiliency Manager, Infrastructure Management Server (IMS), or Replication Gateway</li> <li>■ Define the resiliency domain through Getting Started wizard</li> <li>■ Add assets to your resiliency domain for discovery:                         <ul style="list-style-type: none"> <li>■ Virtual machines</li> <li>■ Applications</li> <li>■ Storage enclosures</li> </ul> </li> </ul>
Discover	<ul style="list-style-type: none"> <li>■ Resiliency Platform's deep discovery enables identification of the following:                         <ul style="list-style-type: none"> <li>■ Virtual machines</li> <li>■ Applications</li> <li>■ Storage enclosures</li> <li>■ Software/hardware replication</li> <li>■ Virtual networks (vSwitches)</li> </ul> </li> </ul>
Define service level objective	<ul style="list-style-type: none"> <li>■ Configure service level objective based on the intended Recovery Point Objective (RPO). service level objective driven configuration enables the following capabilities:                         <ul style="list-style-type: none"> <li>■ Basic monitoring of assets</li> <li>■ Recovery of assets</li> <li>■ Recovery of multi-tier business services (VBS)</li> <li>■ Health status reporting of assets</li> <li>■ Risk alerts and notifications for protected assets</li> </ul> </li> </ul>
Manage	<ul style="list-style-type: none"> <li>■ Single-click rehearsal for resiliency groups and VBS validates disaster readiness:                         <ul style="list-style-type: none"> <li>■ Automated rehearsal</li> <li>■ Automated rehearsal cleanup</li> <li>■ Option of network isolation for workloads during rehearsal</li> </ul> </li> <li>■ Single-click recovery or migration of resiliency groups and VBS:                         <ul style="list-style-type: none"> <li>■ Automated recovery or migration based on the service level objective</li> <li>■ Recovery with predefined network customization</li> <li>■ Recovery based on predefined grouping or order</li> <li>■ Controlled recovery using Resiliency Plans</li> </ul> </li> <li>■ Single-click evacuation plan for resiliency groups and VBS:                         <ul style="list-style-type: none"> <li>■ Option of defining priority levels for VBS</li> <li>■ Automated rehearsal or cleanup rehearsal for evacuation plan</li> </ul> </li> </ul>

# Index

## D

- Data Gateway 14
- deployment
  - replication 13

## O

- overview
  - Data Gateway 11
  - IMS 9
  - Replication Gateways 10
  - resiliency domain 12
  - Resiliency Manager 8

## R

- replication
  - Direct 14
  - Object Storage 14
- Resiliency Platform
  - about 5
  - features and components 6