Veritas™ Resiliency Platform Overview and Planning Guide

VERITAS™

Veritas™ Resiliency Platform Overview and Planning Guide

Last updated: 2019-12-02

Document version: Document version: 3.4 Rev 0

Legal Notice

Copyright © 2019 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/licensing/process

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

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

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

Veritas Technologies LLC

500 E Middlefield Road
Mountain View, CA 94043

http://www.veritas.com

Technical Support

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

https://www.veritas.com/support

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

https://my.veritas.com

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

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

Documentation

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

https://sort.veritas.com/documents

Documentation feedback

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

vrpdocs@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

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 licensable 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 VMware virtual machines
- For recovery of Linux and Windows physical machines to VMware virtual machines

- For recovery of VMware and Hyper-V virtual machines to cloud data center
- For recovery of vCloud virtual machines to vCloud 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 Domain" on page 12.

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 9.

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 9.
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. If you plan to use any third party replication technology, you do not need to deploy Replication Gateway. See "About Replication Gateways" on page 11.
Data Gateway	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 recovery of your data center assets to AWS data center. See "About Data Gateway" on page 11.
Resiliency Domain	The logical scope of a Resiliency Platform deployment that includes Resiliency Manager , IMS, and hosts in both the data centers. It can extend across multiple data centers. See "About Resiliency Domain" on page 12.
Data center	For a disaster recovery use case, the resiliency domain must contain at least two data centers in different locations, a source data center and a target 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. The asset infrastructure can include virtual machines, physical machines (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.
NetBackup Server	The component that allows restoration of virtual machines to a local or remote data center using NetBackup generated backup images. You need to add this component only if you want to restore the virtual machines using NetBackup generated backup images.
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. 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.
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.

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 since the database gets synchronized, all the consoles show the same data. This ensures resiliency and high availability of data in case of any one Resiliency Manager being unavailable. Operations can be performed from any console and the results are displayed 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

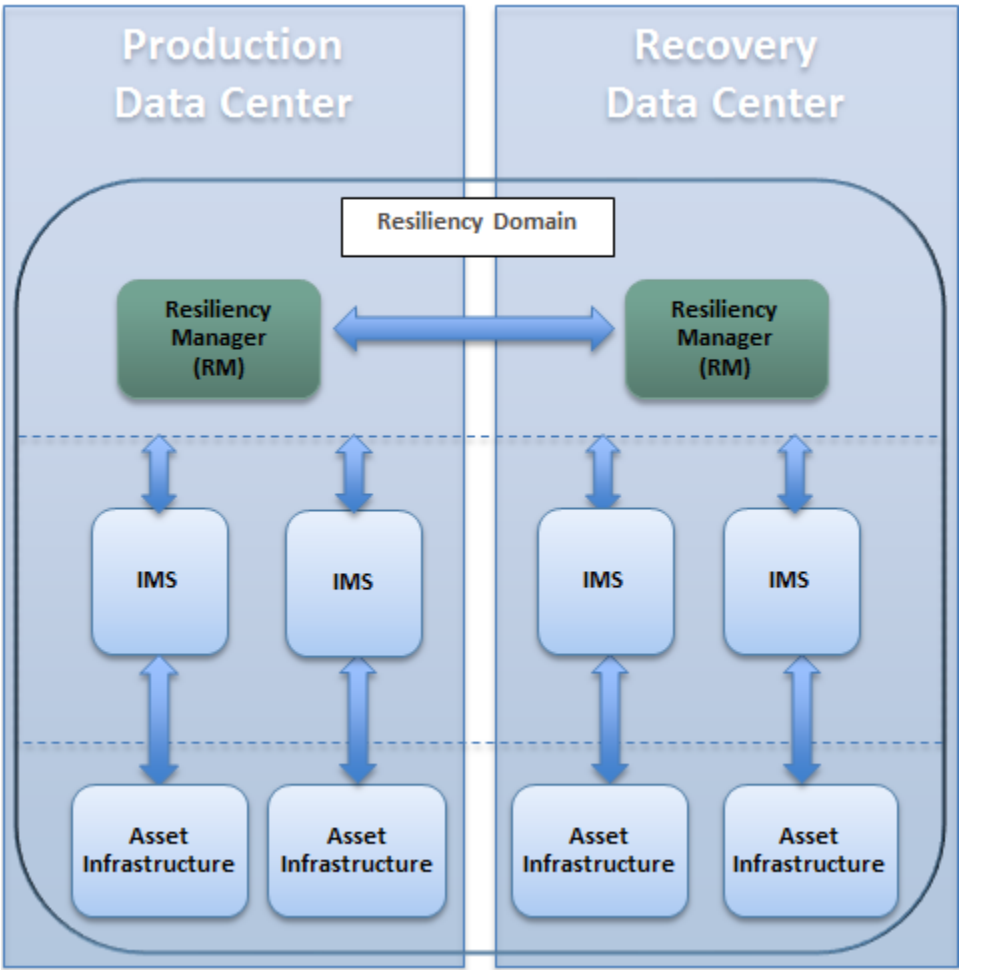

See “About Resiliency Domain” on page 12.

See “About Resiliency Manager” on page 9.

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 source as well as 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 Gateway also performs data optimization like local deduplication and compression. The Gateway on source data center is always paired with a Gateway on target data center. The target data center Gateway is a staging server that applies the data from the source 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

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 such as IMS, along with the infrastructure that is being managed and protected. Within the resiliency domain, Resiliency Platform can protect assets 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 first Resiliency Manager in your environment and access the Resiliency Manager console.

For disaster recovery, the resiliency domain must contain at least two data centers: a source data center and a target 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.

The asset infrastructure includes the data center assets that you add to Resiliency Platform for IMS discovery and monitoring. The asset infrastructure can include virtual and physical hosts (Windows or Linux servers) and virtualization servers for Hyper-V and VMware.

Depending on the technology used for replication, the asset infrastructure can also include Replication Gateways or enclosures (storage arrays). This diagram does not show the replication details for the asset infrastructure.

Figure 1-2 Resiliency Platform components in resiliency domain

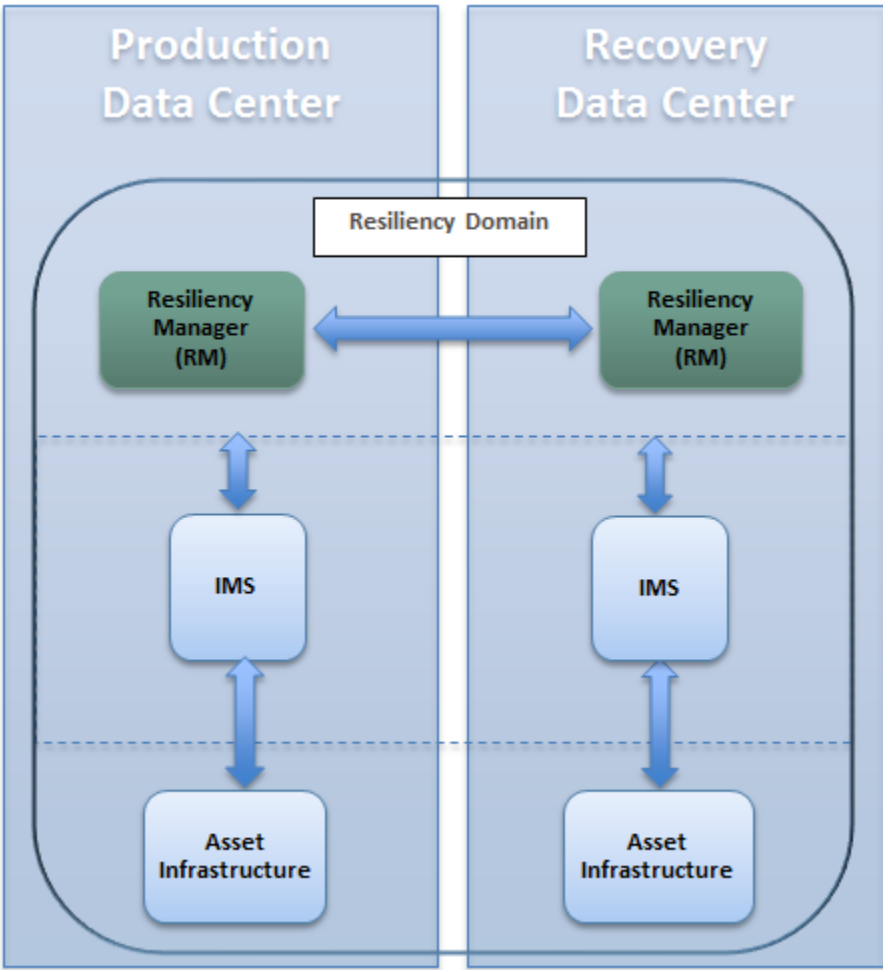

See “Replication in a Resiliency Platform deployment” on page 23.

See “About Resiliency Manager” on page 9.

See “About Infrastructure Management Server (IMS)” on page 9.

See “Planning a resiliency domain for efficiency and fault tolerance” on page 25.

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 1-1

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
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 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

Chapter 2

Planning your environment for disaster recovery using Resiliency Platform

This chapter includes the following topics:

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. For more information on VMware API I/O filtering framework, refer to VMware documentation.

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.

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 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.

How Veritas Resiliency Platform Data Mover handles virtual machine writes

Veritas Resiliency Platform Data Mover uses the vtstap service running on ESXi 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 a lag between the data on 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 virtual machine 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 transceiver on the source Replication Gateway sends the I/Os across the network to the transceiver on the target Replication Gateway.
- The I/Os are then sent to the applier and the applier writes them to the target data center storage.

Replicating thin disks

Resiliency Platform supports protecting the virtual machines with thin disks optimally by provisioning corresponding thin disks on the target data center. Resiliency Platform Data Mover replicates only the used blocks from the disks on the source data center to the disks on the target data center thereby maintaining the thin nature of the disks. In addition to thin disks, Resiliency Platform Data Mover also replicates only the used blocks from a thick (lazy zero provisioned) disks to the disks on the target data center.

If you have NFS based datastores on the source data center, then irrespective of the disk type on the source data center, a thick disk is created on the target data center. This is required as NFS layer is not thin-aware.

When you upgrade from 3.1 or an earlier version to version 3.2 or later and if you have thin disks on the source data center and the resiliency group has been configured for disaster recovery (DR), you need to delete the existing resiliency groups after upgrading. Reconfigure the resiliency groups to utilize the thin provisioning feature. If you are upgrading from 3.1 Update 1, you do not need this step.

Architecture of Resiliency Platform Data Mover

When you want to recover VMware virtual machines from an on-premises data center to an on-premises data center then the replication technology used is Resiliency Platform Data Mover with VMware vSphere APIs for I/O Filtering (VAIO) framework. The framework provides an I/O filter for replication that has been certified by VMware. The filter runs inside an ESXi server and intercepts I/O requests moving between a guest operating system and virtual disks. For more information on VMware API I/O filtering framework, refer to VMware documentation.

Resiliency Platform Data Mover deploys a user-land module, called **vtstap**, on the ESXi host where the protected virtual machines are running. This module is built using the VMware VAIO APIs to intercept and replicate I/Os from the virtual machines.

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.
- 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 takeover) is run. 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. This establishes the replication channel between the source and the target data centers. You can choose to encrypt the communication between gateways, unless you are using a dedicated VPN link.

Figure 2-1 Replication architecture

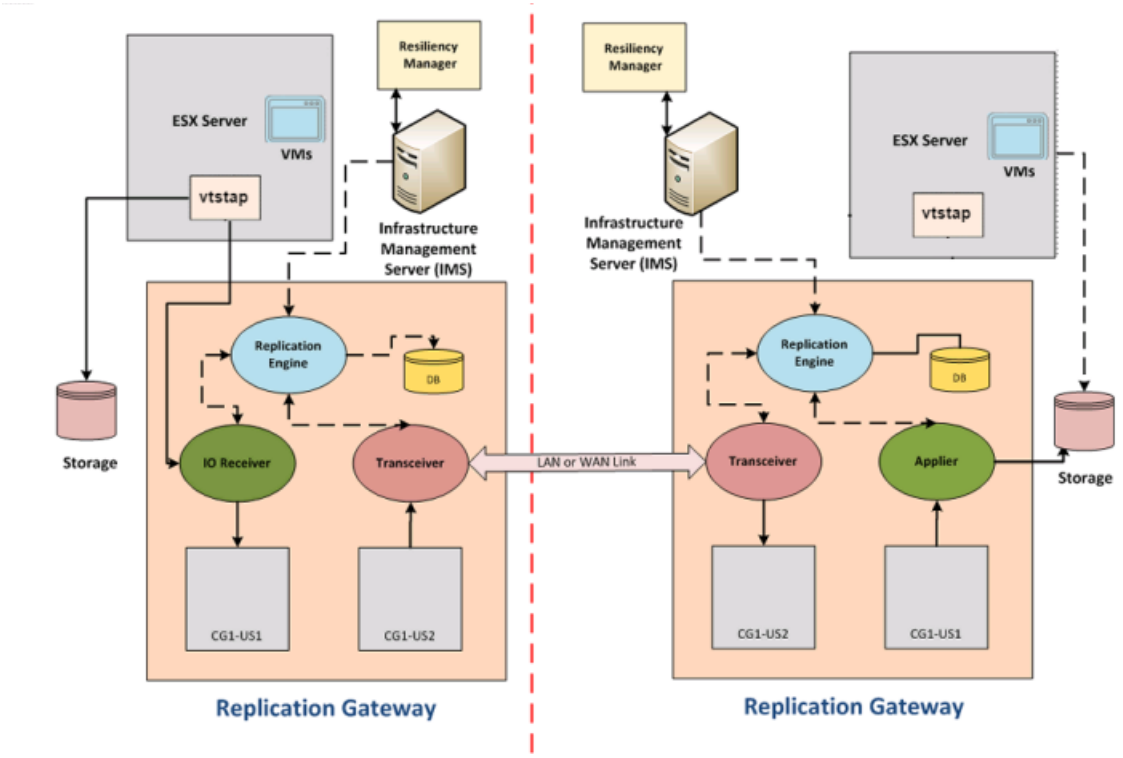

About synchronization used by Resiliency Platform Data Mover

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 21.
- See "About incremental synchronization" on page 22.

You do not have an option to choose between the two types of synchronization techniques. The synchronization technique being used at any time is decided internally by Resiliency Platform based upon the requirements.

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 data center to provide business continuity. After a takeover, the virtual machine runs in the target data center. Once the source data center is back up and running, you must perform a Resync operation from the target data center before you can migrate back to the source data center. This Resync operation launches a full synchronization to synchronize the data on the source data center with the data in the target data center. When the synchronization completes, the source 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 used by Resiliency Platform Data Mover” on page 21.

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 used by Resiliency Platform Data Mover" on page 21.

Replication in a Resiliency Platform deployment

Veritas Resiliency Platform supports several forms of replication for data recovery from your source data center to your target 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 2-2 Replication in a Resiliency Platform deployment

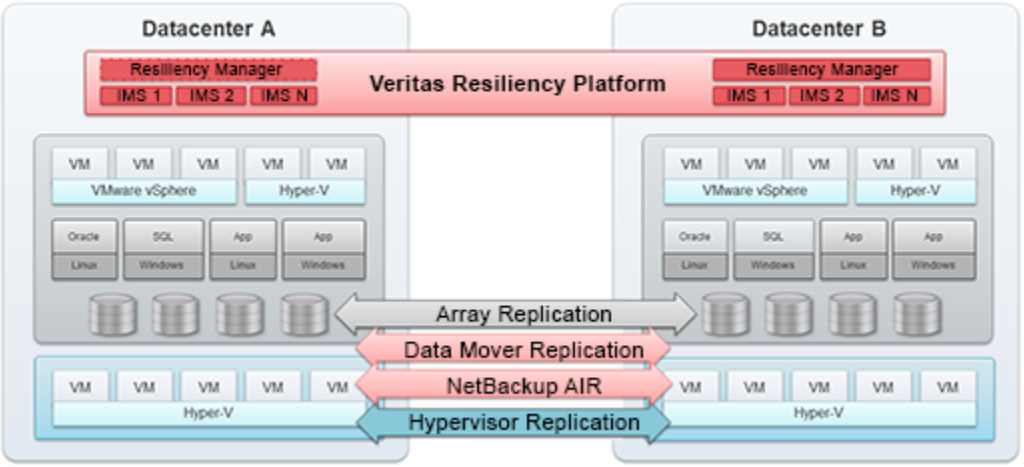

When you use Resiliency Platform Data Mover for recovery of your data center assets to AWS, you have an option to choose between the two modes of replication: Direct mode and Object Storage mode.

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 source data center is always paired with a Replication Gateway on target data center. In Direct mode replication, the Replication Gateway at source data center directly communicates with the Replication Gateway at target data center. The source site Replication Gateway acts as a staging server that aggregates and batches data from multiple virtual machines during replication. The target 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. To enable

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

See "About Data Gateway" on page 11.

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 on the target data center does not hamper the replication.
  In case where the Replication Gateway at target data center goes down, data synchronization is performed after the Gateway is restored back or replaced. If the Replication Gateway at the source data center goes down before the synchronization gets completed, you may loose some of the data if you are using the Direct mode of Replication. But if you are using the Object Storage mode replication, then the data can be pulled from S3 bucket and you do not loose any data.

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.

If you have multiple Resiliency Managers in the resiliency domain, the recommended WAN latency between two Resiliency Managers is recommended to be less than 30 milliseconds. The maximum WAN latency allowed between the two Resiliency Managers is 50 milliseconds. This latency requirement applies to any two Resiliency Managers - within a single data center or across two different data centers.

The recommended minimum deployment for disaster recovery to on-premises data center in production environment would be four virtual appliances: a Resiliency Manager and Infrastructure Management Servers (IMS) in the source data center and a Resiliency Manager and IMS in the target data center. In a test environment, you can have only one Resiliency Manager at the target data center.

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

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

Both the source and the 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. The minimum size of the external storage must be 50GB.

Veritas Resiliency Platform is dependent on host name resolution to work between the resiliency platform appliances and its configured assets. Trying disaster recovery when there is a partial network infrastructure failure, such as DNS failure, can lead to incomplete disaster recovery activities, if the appliances and assets are not configured using network host names.

In order to avoid encountering such issues, you must use fully qualified network host names in the following scenarios:

- Joining a new resiliency manager to a domain
- Adding an infrastructure management server to a data center
- Adding a resiliency platform datamover to a data center
- Preparing hosts for replication using resiliency platform data mover
- Adding application hosts
- Adding hyper-v servers

Table 2-2 Reference topics

For more information on...	Refer to...
Network	About FIPS enablement About NAT support About IPv6 support
Capacity planning for Resiliency Platform replication appliances	VMware vSphere environment VMware vSphere and Azure environment VMware vSphere and AWS environment
Network objects and mapping of network objects required across data centers	About network objects Creating network pairs between data centers

See "About Resiliency Domain" on page 12.

See "About Resiliency Manager" on page 9.

See "About Infrastructure Management Server (IMS)" on page 9.

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	■ 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

Table 2-3 On-boarding with Resiliency Platform *(continued)*

Step	Description
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
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

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, virtualization servers, virtual machines, enclosures, and applications.
assets	The virtual machines, physical machines, or applications that have been discovered by the Infrastructure Management Server (IMS) and that can be grouped into resiliency groups.
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 source data center and target data center. Each data center has a Resiliency Manager and one or more IMSs.
host	In Veritas Resiliency Platform, the term hosts means Application host, Resiliency Platform Data Mover host, Storage discovery host, VMware Discovery host, and Hyper-V host.
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.
klish	Command Line Interface SHell. Provides the command line menu on the virtual appliance for use after the initial bootstrap configuration.
migrate	A planned activity involving graceful shutdown of assets at the source data center and starting them at the target data center. In this process, replication ensures that consistent data is made available at the target 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.
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.
Replication Gateway	The Veritas Resiliency Platform component that performs data replication between the source and the target data center.
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 to be 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 web 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.
Resiliency Platform Data Mover Replication host	To enable replication using Resiliency Platform Data Mover replication technology, you need to add an asset and prepare it for replication. Asset can be a physical machine or a virtual machine.
source data center	The data center that is normally used for business.
take over	An activity initiated by a user when the source data center is down due to a disaster and the assets need to be restored at the target data center to provide business continuity.
target data center	The data center that is used if a disaster scenario occurs.
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 operations are performed on the resiliency groups.
VAIO framework	VMware framework consisting of vSphere APIs for I/O Filtering. This framework enables Veritas Resiliency Platform to run filters on ESXi servers and intercept any I/O requests from a guest operating system to a virtual disk.
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 recovery in case of a disaster in the desired order.
Veritas Replication Set	A virtual machine, which belongs to the resiliency group, is termed as Veritas Replication Set. All the disks attached to this virtual machine, including the boot and data disk, constitute a Veritas Replication Set. The write order fidelity is maintained across all disks in a given replication set.
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.