



Hewlett Packard
Enterprise

HPE 3PAR Remote Copy Concepts Guide

Abstract

This guide is for all levels of system and storage administrators who plan, build, and manage HPE 3PAR Remote Copy system configurations. This guide explains the concepts, terminology, and fundamentals of Remote Copy.

Part Number: QL226-10537
Published: September 2018
Edition: 1

Notices

The information contained herein is subject to change without notice. The only warranties for Hewlett Packard Enterprise products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. Hewlett Packard Enterprise shall not be liable for technical or editorial errors or omissions contained herein.

Confidential computer software. Valid license from Hewlett Packard Enterprise required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Links to third-party websites take you outside the Hewlett Packard Enterprise website. Hewlett Packard Enterprise has no control over and is not responsible for information outside the Hewlett Packard Enterprise website.

Acknowledgments

VMware[®], VMware[®] vCenter[™], and VMware vSphere[®] are registered trademarks or trademarks of VMware, Inc. in the United States and/or other jurisdictions.

Contents

- HPE 3PAR Remote Copy overview..... 5**

- HPE 3PAR Remote Copy features.....6**

- Remote Copy terminology.....7**
 - Remote Copy groups..... 8
 - Remote Copy links..... 8
 - Remote Copy targets..... 9
 - Remote Copy target definitions..... 9
 - Recovery Point Objective..... 10
 - Recovery Time Objective..... 11

- HPE 3PAR Remote Copy options and requirements..... 12**

- Network transport methods for Remote Copy.....13**
 - Remote Copy over IP (RCIP) network transport requirements..... 13
 - Remote Copy over Fibre Channel (RCFC) network transport requirements..... 14

- Replication directions in Remote Copy.....15**
 - Direction of data flow..... 15

- Replication modes for Remote Copy.....16**
 - Synchronous mode for Remote Copy..... 16
 - Advantages of synchronous mode replication..... 17
 - Disadvantages of synchronous mode replication..... 17
 - Asynchronous periodic mode for Remote Copy..... 17
 - Advantages of asynchronous periodic mode replication..... 18
 - Disadvantages of asynchronous periodic mode replication 18
 - Asynchronous streaming mode for Remote Copy..... 18
 - Advantages of asynchronous streaming mode replication..... 19
 - Disadvantages of asynchronous streaming mode replication..... 19

- Configuration options for Remote Copy.....20**
 - 1-to-1 unidirectional Remote Copy configuration..... 20
 - 1-to-1 bidirectional Remote Copy configuration..... 21
 - M-to-N Remote Copy configurations..... 21
 - Synchronous long-distance Remote Copy configuration..... 22
 - Remote Copy with HPE 3PAR Peer Persistence configuration..... 23
 - Three Data Center Peer Persistence configuration..... 24

- Synchronization with Remote Copy..... 26**

Types of synchronization for Remote Copy.....	26
Remote Copy volumes and synchronization.....	26
Snapshots and coordinated snapshots within Remote Copy.....	27
Synchronization with synchronous mode volume groups.....	27
Synchronization with asynchronous periodic mode volume groups.....	27
Automatic resynchronization for asynchronous periodic mode volume groups....	28
Synchronization with asynchronous streaming mode.....	28
Snapshot frequency for asynchronous streaming mode volume groups	28
Resynchronization order for asynchronous streaming mode volume groups.....	29

System access, requirements, and best practices for Remote Copy..... 30

System access to set up Remote Copy.....	30
Remote Copy requirements and best practices.....	30
Prerequisites for setting up Remote Copy.....	30
Software restrictions for Remote Copy.....	31
Best practices for Remote Copy.....	31
Best Practice for Remote Copy volumes.....	31

Documentation, support, and other resources..... 33

Supporting documentation for Remote Copy.....34

Support and other resources..... 35

Accessing Hewlett Packard Enterprise Support.....	35
Accessing updates.....	35
Customer self repair.....	36
Remote support.....	36
Warranty information.....	36
Regulatory information.....	37
Documentation feedback.....	37

HPE 3PAR Remote Copy overview

HPE 3PAR Remote Copy software is a data replication solution for HPE 3PAR StoreServ Storage systems. Data can be replicated to remote locations which allows the protection and sharing of data. Remote Copy when combined with appropriate clustering technology, services, and processes, can serve as the foundation for a Disaster Recovery solution.

This part describes HPE 3PAR Remote Copy features and introduces you to the terms used.

HPE 3PAR Remote Copy features

Remote Copy features and benefits include:

- **Set up and test basic configurations quickly** using HPE 3PAR StoreServ Management Console or HPE 3PAR Command-Line Interface.
- **Leverage existing infrastructure** with native IP and Fibre Channel support eliminating the need for dedicated or specialized infrastructure.
- **Provide flexible and efficient multitenant disaster recovery** with different configuration options and licensing for only a portion of installed capacity.
- **Achieve fast recovery time objectives (RTO) and stringent recovery point objectives (RPO)** with multisite, multinode synchronous long-distance replication configuration.
- **Replicate only written data for bandwidth optimization** improves the speed of initial synchronization and enables conversion of fully provisioned volumes to thin-provisioned volumes.
- **Automate transparent failover** between sites at metropolitan distances, with an option to have a third copy at a geographic distance with HPE Peer Persistence.
- **Synchronous replication** between sites at metropolitan distances, with the option to have a third copy at a geographic distance site with synchronous long distance (SLD).

Remote Copy terminology

Remote Copy group

A Remote Copy virtual volume set of one or more virtual volumes to be replicated to another system. I/O consistency is guaranteed across the virtual volumes in a Remote Copy group.

Primary system

The HPE 3PAR StoreServ Storage system which contains the Remote Copy group *from* which data is replicated to a secondary system. The primary system takes write I/Os from a server for the virtual volumes in a Remote Copy group and *sends* that data to a secondary system.

A primary system can also be referred to as the source or local system.

Secondary system

The storage system which contains the Remote Copy group *to* which data is replicated. The secondary system receives replication data from the primary system and *applies* that data to the virtual volumes in a Remote Copy group.

The secondary system can also be referred to as the target or remote system. Target is often used in other documentation to represent the secondary system.

Primary-rev

Following a failover, the secondary system takes on the role of primary-rev indicating that it is now the primary storage system as the result of a failover.

Secondary-rev

Following a failover, the primary system takes on the role of secondary-rev indicating that it is now the secondary storage system as the result of a failover.

Source system

The system that generates the data and sends it to the target. Because the primary and secondary roles can change, Remote Copy uses the term **source system** to refer to a storage system in a linked pair of storage systems. The source system is the primary or primary-rev system.

Target system

The system that receives the replicated data. Because the primary and secondary roles can change, Remote Copy uses the term **target system** to refer to the other storage system in a linked pair of storage systems.

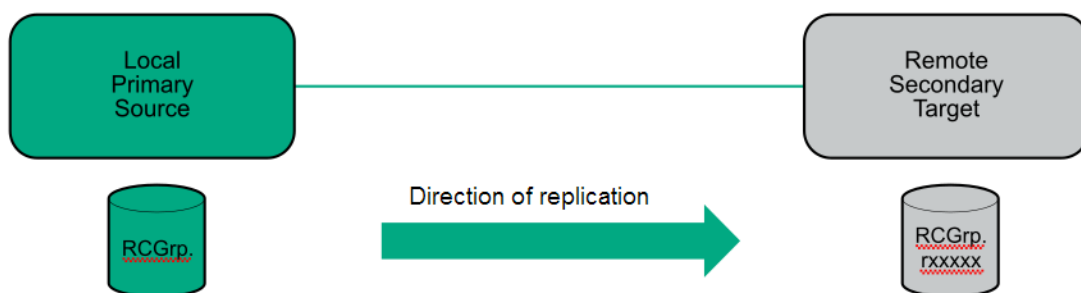


Figure 1: Direction of replication example




Figure 2: Primary failure example

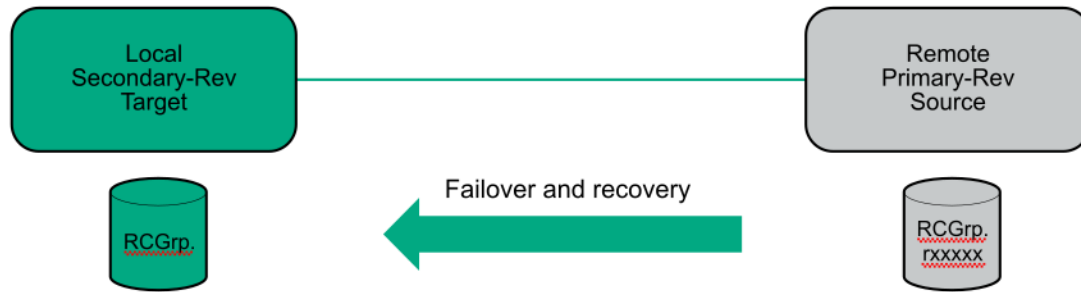


Figure 3: Failover and recovery example

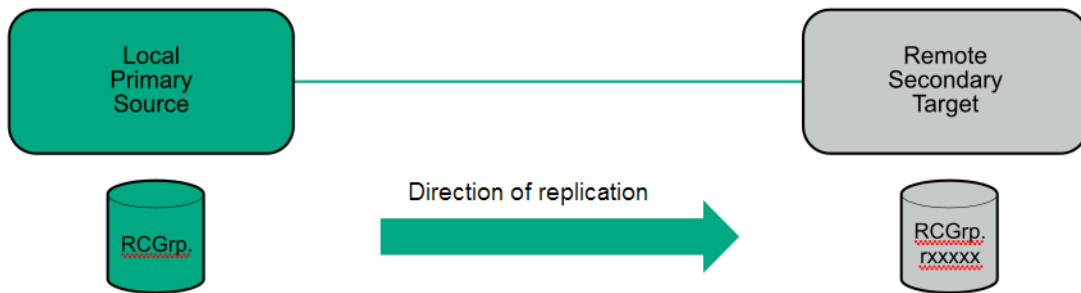


Figure 4: Restoration example

In a failover, the original direction of replication is restored after recover and restore actions are completed. In a reverse, the original direction of replication is restored when the primary-rev system becomes the primary and the secondary-rev system becomes the secondary.

Remote Copy groups

A Remote Copy group is a group of one or more virtual volumes to be replicated to another system.

Because the volumes in a Remote Copy group are related, Remote Copy ensures that the data in the virtual volumes within the group maintain write consistency.

Remote Copy operations are performed on the Remote Copy group rather than individual volumes.

For example, when you start or stop Remote Copy, operations are started and stopped for the whole Remote Copy group.

Remote Copy links

Remote Copy links establish communication between storage systems.

When you create target definitions for a system during setup, define a minimum of one pair of links for each secondary system. For redundancy, configure more than one communication network between storage systems. Replication can then continue even if one of the networks fails.

You create links during Remote Copy setup by using either the `creatercopytarget` command or the `admitrcopylink` command.

When you create a link, it becomes both a sending and receiving link for the port specified in the command.

Remote Copy targets

When you set up HPE 3PAR Remote Copy, you configure the secondary system as the target of the primary system, and the primary system as the target of the secondary system.

For example, in a Remote Copy configuration composed of `System1` and `System2`, the target system for `System1` is `System2`, and the target system for `System2` is `System1`.

In Remote Copy, "target" has a special meaning. A primary or primary-rev system sends data for a Remote Copy group using a particular target on a secondary system. The "target" specified indicates which physical links and replication mode are used to send the data. A target can have only one replication mode. Creating targets on both systems in the Remote Copy configuration:

- Links the systems.
- Allows for data flow in either direction, depending on the configuration and on temporary situations such as downtime or failure and disaster recovery.

Remote Copy target definitions

A target definition identifies the destination for data being replicated during Remote Copy operations. The secondary (or secondary-rev) system is the other system in the linked set of storage systems.

You specify target definitions during Remote Copy configuration.

For example, in a `System1/System2` Remote Copy configuration, you specify target definitions as follows:

- On `System1`, the target definition is `System2`.
- On `System2`, the target definition is `System1`.

The name used for the target definition is not required to match the system name.

This relationship between target definitions and the linked storage systems is the same for all Remote Copy configurations.

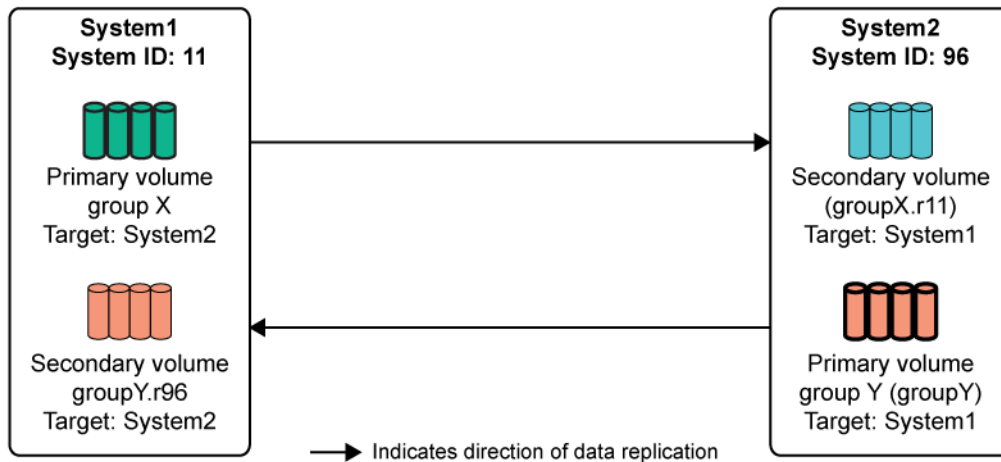


Figure 5: Target definitions for a bidirectional Remote Copy configuration - example

The target definitions used by each Remote Copy group in the example are:

Volume group and location	Target definition/secondary system
For primary volume group <code>groupX</code> on <code>System1</code>	<code>System2</code>
For secondary volume group <code>groupX.r11</code> on <code>System2</code>	<code>System1</code>
For primary volume group <code>groupY</code> on <code>System2</code>	<code>System1</code>
For secondary volume group <code>groupY.r96</code> on <code>System1</code>	<code>System2</code>

Recovery Point Objective

Recovery Point Objective (RPO) is a measure of how much data can be lost were a disaster to occur. RPO refers to the point in time in the past to which you will recover.

When designing a replication solution, the RPO generally drives the replication mode chosen and the sizing of the links used for replication. If the primary system fails and a failover occurs, the RPO is a worst case measurement.

RPO is measured in a quantity of time and not in a quantity of data. For example:

- RPO = 0 means there is no data loss as a result of a disaster. It is the best result with synchronous replication.
- RPO = 2 minutes is the best result with asynchronous streaming replication.
- RPO = 10 minutes is the best result with asynchronous periodic replication. For an RPO over 10 minutes, use asynchronous periodic replication.

For example, a disaster hits at 2:00 pm. After recovery operations at the disaster recovery site complete, an RPO of two hours would guarantee that all transactions that were committed up to and including 12:00 pm would be present in the replicated copy of a database. The RPO could, in fact, contain transactions committed after 12:00 pm but the defined two-hour RPO guarantees transactions committed up to 12:00 pm will always be present.

Recovery Time Objective

The Recovery Time Objective (RTO) is the maximum amount of downtime that elapses following a disaster before the application is back online and processing new transactions. RTO refers to the point in time in the future at which you will be up and running again. An RTO is measured in seconds, minutes, hours, or days.

In most disaster-tolerant designs, RTO is secondary to the amount of data loss that can be tolerated.

For example, a disaster hits at 2:00 pm. If the RTO is 1 hour, your expectation is that the application will be back online and ready to process transactions by 3:00 pm.

HPE 3PAR Remote Copy options and requirements

This part describes the various options available and requirements necessary to set up your Remote Copy system. Replication mode, network transport method, and configuration type are some of the factors in deciding how you want to set up your Remote Copy system.

Network transport methods for Remote Copy

Storage systems in a Remote Copy configuration are connected through a dedicated link or through a network.




Figure 6: Remote Copy links

Remote Copy supports the following network transports:

- **IP networks**
Remote Copy over IP (RCIP) links two systems over Ethernet ports.
- **Fibre Channel networks**
Remote Copy over Fibre Channel (RCFC) links systems using native Fibre Channel ports.

HPE 3PAR Remote Copy does not support network address translation (NAT) over RCIP.

Remote Copy over IP (RCIP) network transport requirements

For RCIP configurations, each link between a pair of storage systems is a logical link between a controller node on one storage system and a controller node on the other storage system in the configuration. These links use a network adapter (NIC) port from each of the nodes in the storage systems.

To use an IP network for Remote Copy, the following requirements must be met.

- A minimum of two links (a node pair) for each storage system in the Remote Copy configuration.
Utilizing like node pairs (node pair: 0/1 or 2/3) on each storage system is recommended, but is not a requirement.
- The NIC interface and the management Ethernet port of the HPE 3PAR StoreServ Storage system controller node must be on different IP subnets.
- A unique IP address is used for each RCIP NIC interface.
- The Remote Copy subnet can be shared with other I/O traffic, but a guaranteed amount of bandwidth must be dedicated to Remote Copy. The RCIP network must be sized appropriately to replicate the data being generated. Without a guaranteed amount of bandwidth being dedicated to RCIP, the designated RPO for the solution cannot be guaranteed.

When replicating synchronously over RCIP, the bandwidth must be large enough to ensure that no performance issues arise from replicating data synchronously. For asynchronous periodic or

asynchronous streaming replication, the bandwidth must be large enough to ensure that the RPO can be met.

More information

- *HPE 3PAR StoreServ 3PAR OS 3.3.1 Support Matrix*
- *Disaster-tolerant solutions with HPE 3PAR Remote Copy white paper*

Remote Copy over Fibre Channel (RCFC) network transport requirements

For RCFC configurations, each link between storage systems is a physical link between a controller node on one storage system and a controller node on the other storage system in the configuration.

To use a Fibre Channel network for Remote Copy, the following requirements must be met.

- Utilizing like node pairs (node 0 and 1 is an example of a node pair) on each storage system is recommended, but is not a requirement.
- FC SAN is used to connect the storage systems (direct connect FC is not supported).
- Each pair of RCFC ports between the primary and secondary system must be in their own zone. Fabric zones cannot be shared. A single initiator to a single port configuration and a dual fabric layout is preferred.

For example, only two ports per zone are allowed as shown in **Figure 7: Zoning RCFC node pairs** on page 14. Each zone must contain one port from the primary site and a second port from the secondary site.

- Requirements in the *SAN design Reference Guide* are met.
- Interrupt coalescing must be disabled on the RCFC ports.

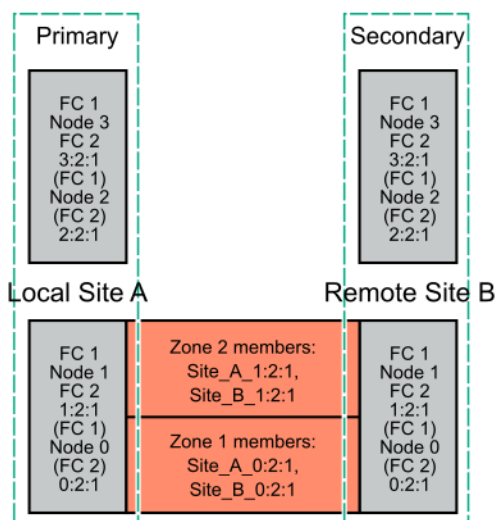


Figure 7: Zoning RCFC node pairs

Replication directions in Remote Copy

Unidirectional – If replication occurs in only one direction in a storage system configuration, the configuration is under unidirectional replication.

Bidirectional – If replication is occurring in both directions in a storage system configuration, the configuration is under bidirectional replication. The systems in the Remote Copy configuration serve as both primary and secondary systems. The system role changes, depending on the volume group. The primary system is always the system with the primary volume groups.

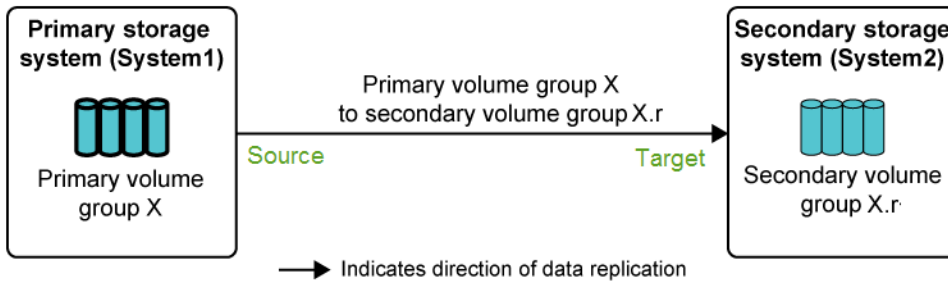


Figure 8: Unidirectional replication example

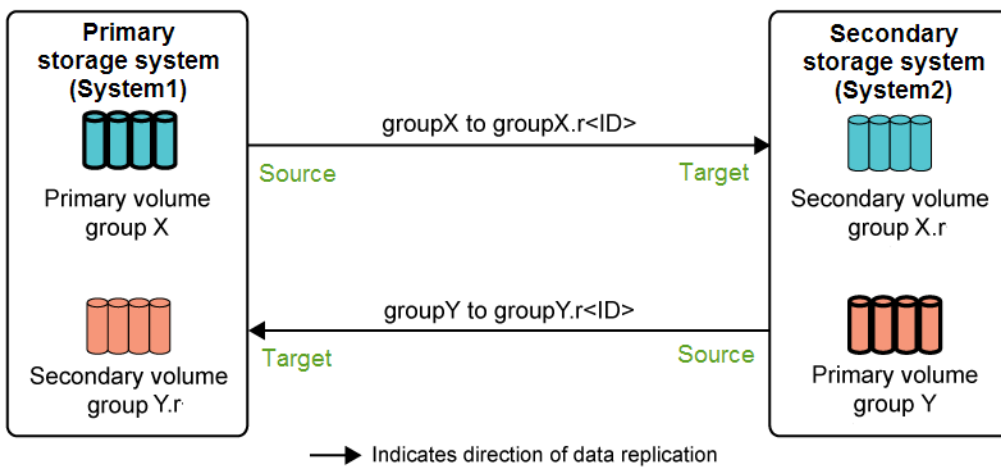


Figure 9: Bidirectional replication example

Direction of data flow

Under normal operating conditions, the natural direction of data flow is from the primary volume group to the secondary volume group. Remote Copy replicates data from primary volume groups on the primary system to secondary volume groups on the secondary system.

Under the following circumstances, temporarily switch the direction of data flow:

- Disaster recovery
- System maintenance/replacement
- Scheduled downtime

Switching the direction of data flow reverses the replication direction of the volume groups.

Replication modes for Remote Copy

The replication modes for Remote Copy volume groups are:

Replication mode	Difference between other modes	RPO
Synchronous (sync): Synchronous mode volume groups stay synchronized as long as the replication links remain up and available.	Data concurrency is greater than in asynchronous periodic mode or asynchronous streaming mode, but host write I/O response times may be higher than in asynchronous periodic or asynchronous streaming mode.	RPO is 0.
Asynchronous periodic (periodic): Asynchronous periodic mode volume groups resynchronize at specific intervals using snapshots.	Host I/O write response times can be better than in synchronous mode, but data is not always in sync.	For properly sized replication links, the RPO is 2 times the replication period chosen for the Remote Copy group (minimum of 10 minutes).
Asynchronous streaming (async): Asynchronous streaming mode volume groups replicate data to the secondary system constantly.	Data concurrency is greater than asynchronous periodic mode, but not as good as synchronous mode. Host I/O write response times can be better than synchronous mode.	For properly sized replication links, the maximum RPO is 2 minutes.

Synchronous mode for Remote Copy

When Remote Copy volume groups operate in synchronous mode, a host write is committed to both the primary and the secondary storage systems before the primary system acknowledges the host write to the server.

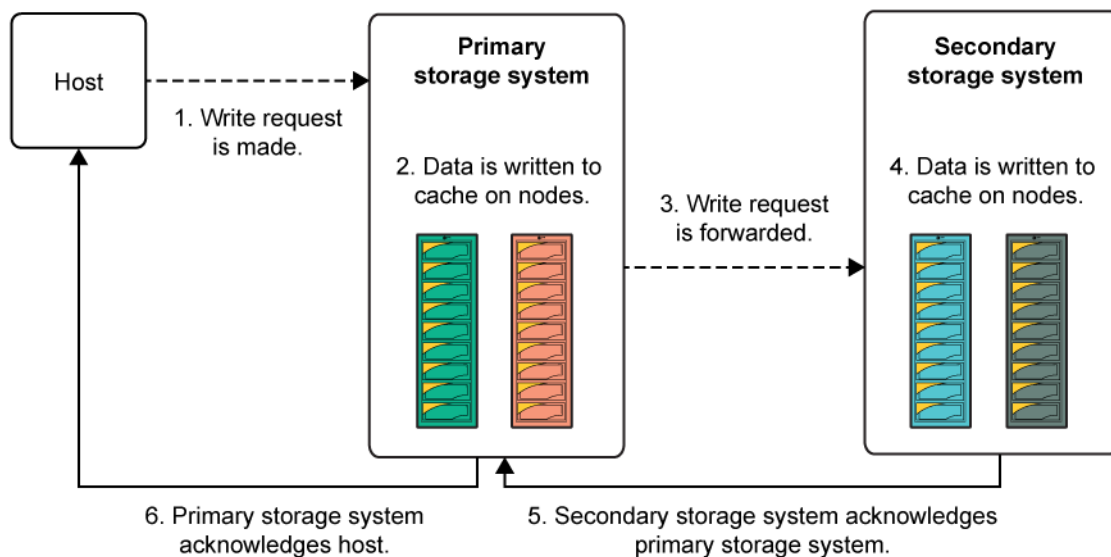


Figure 10: Remote Copy in synchronous mode

Advantages of synchronous mode replication

The advantages of synchronous mode replication is an RPO of zero (no data loss).

- Even if the primary system, the secondary system, or the communication links go down, both storage systems contain all the data for I/O that has already been acknowledged to the server to the point in time that the failure occurred.
- If a disaster affects the primary system and it goes down, the copy of the data that is stored on the secondary system is an exact replica of the data on the primary system and therefore can immediately be used to continue working in the application.

Disadvantages of synchronous mode replication

Synchronous mode replication can impact server performance due to latency issues and link constraints. Consider the following:

- The write I/O service times on the primary and secondary systems
- The time to move the data across the link is a function of the link latency, link speed, and the number of buffers available on the link

As the distance between the primary and secondary storage systems increases, the latency also increases. For example, a one-way distance of 100 kilometers over dark fiber adds approximately 1 ms to the write latency.

Even with the primary and secondary storage systems located side-by-side and linked by an FC SAN, using synchronous mode volume groups adds latency to a host write because the write must be committed on both systems before the IO complete is sent to the server.

Asynchronous periodic mode for Remote Copy

When Remote Copy volume groups operate in asynchronous periodic mode, the secondary system is periodically updated with the write data issued by the server as of the last resynchronization update. The write is acknowledged as soon as data resides in the cache of the secondary system. The synchronization period is set up during configuration of the Remote Copy group.

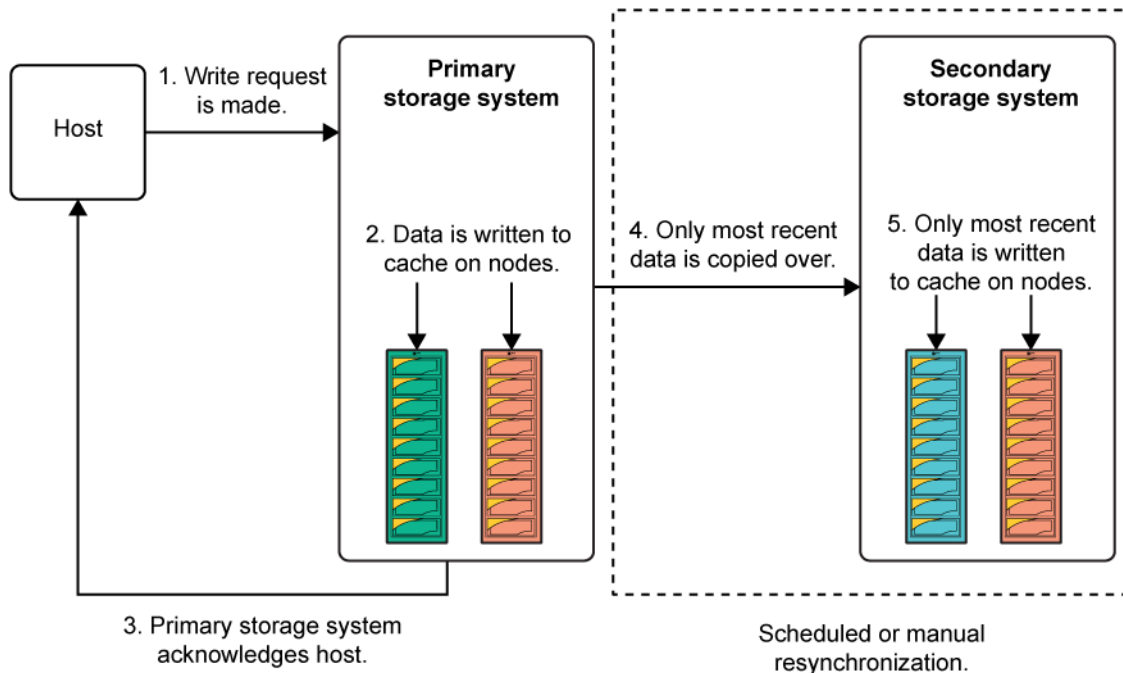


Figure 11: Remote Copy in asynchronous periodic mode

Advantages of asynchronous periodic mode replication

Remote Copy has no negative effect on host write I/O latency when used with replication links that have high latency. If one area of a volume is written to multiple times between two resynchronizations, Remote Copy sends only the last write to the secondary system. Therefore, asynchronous periodic mode Remote Copy may transfer less data overall than synchronous mode Remote Copy.

Disadvantages of asynchronous periodic mode replication

The disadvantage of asynchronous periodic mode replication is:

A higher RPO. The RPO is two times the resynchronization interval, the time set for the synchronization period.

For example, if you set the synchronization period to 5 minutes, Remote Copy queues up the data for 5 minutes, and then sends the data for 5 minutes. The RPO=10 minutes.

More information

Disaster-tolerant solutions with HPE 3PAR Remote Copy white paper

Asynchronous streaming mode for Remote Copy

When Remote Copy volume groups operate in an asynchronous streaming mode, asynchronous I/O operations are kept in cache memory until they are written on the secondary system. Then an acknowledgment is sent back to the primary system. The asynchronous I/O operations are written in the secondary volumes in the same order they were written in the primary volumes.

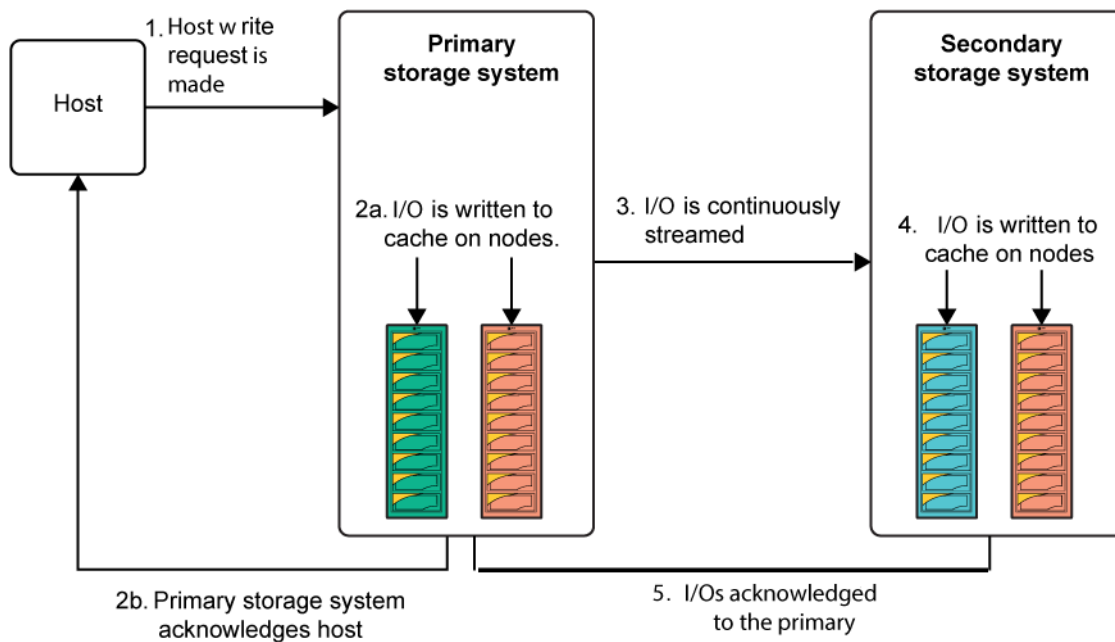


Figure 12: Remote Copy in asynchronous streaming mode

Advantages of asynchronous streaming mode replication

The advantages of asynchronous streaming mode replication are:

- Provides higher replication latency support and small RPOs without impacting server I/O performance like synchronous replication.
- Delivers an RPO in the range of seconds compared to asynchronous periodic replication.
- Host write I/O latency is not affected by the network latent, latency jitter, or packet loss on the network.

Disadvantages of asynchronous streaming mode replication

The disadvantages of asynchronous streaming mode replication are:

- Data consistency is ensured at the secondary storage system, but the data may not be current (RPO > 0).
- Network bandwidth should not be shared with other modes of Remote Copy replication.

Configuration options for Remote Copy

Remote Copy configurations are based on the relationship between the storage systems in the configuration. Within a Remote Copy configuration, the primary storage system is the system that holds the volumes that are copied to the secondary storage system.

Remote Copy supports the following configurations:

Basic solution types

- **One-to-One (1-to-1):** One pair of storage systems that replicate data between one another.
 - **1-to-1 unidirectional Remote Copy configuration** on page 20
 - **1-to-1 bidirectional Remote Copy configuration** on page 21

Advanced solution types

- **M-to-N Remote Copy configurations on page 21:** Up to five storage systems that replicate their data among one another.
- **Synchronous long-distance Remote Copy configuration on page 22:** A storage system replicates its data to two target systems. The same Remote Copy volume group is replicated to the two target systems, one in synchronous mode, and the other in asynchronous periodic mode.
- **Remote Copy with HPE 3PAR Peer Persistence configuration on page 23:** A high availability (HA) solution that leverages synchronous mode Remote Copy and ensures that the same data is on both arrays. Both source and target volumes are exported to the hosts in the Peer Persistence configuration. Peer Persistence offers automated transparent failover in the event the primary system suffers a failure.
- **Three Data Center Peer Persistence configuration on page 24:** 3DC-PP is a Peer Persistence configuration where the data is also replicated to a third system. This solution provides disaster recovery for three data centers at different locations.
- **HPE Cluster Extension (CLX):** A HA solution that leverages Remote Copy and offers automatic failover over distances. See *HPE 3PAR StoreServ Cluster Extension Software Administrator Guide* at the HPE Support Center <http://www.hpe.com/support/hpesc>.
- **HPE 3PAR Storage Replication Adapter (SRA) Software for VMware Site Recovery Manager (SRM):** HPE 3PAR Storage Replication Adapter (SRA) is a plug-in to VMware vCenter Site Recovery Manager. 3PAR SRA can execute specific Remote Copy functions to build, manage, test, and execute disaster recovery for virtualized systems. See *HPE 3PAR Storage Replication Adapter for VMware® vCenter Site Recovery Manager User Guide* at the HPE Support Center <http://www.hpe.com/support/hpesc>.

1-to-1 unidirectional Remote Copy configuration

In a 1-to-1 unidirectional Remote Copy configuration, each storage system in the pair plays only one role: one system is the primary system, and one system is the secondary system.

The primary storage system holds only primary volume groups, and the secondary storage system holds only secondary volume groups.

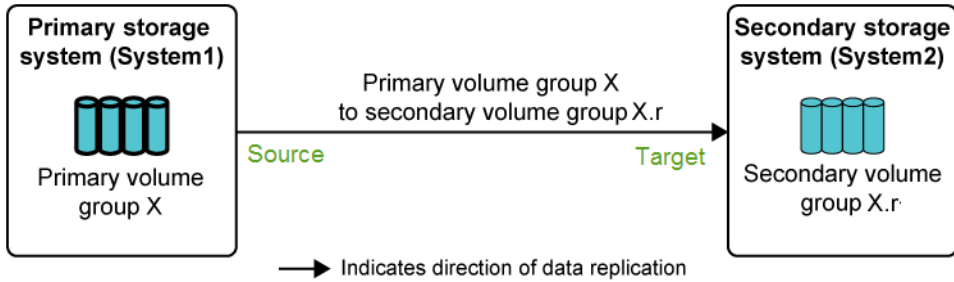


Figure 13: 1-to-1 unidirectional Remote Copy

1-to-1 bidirectional Remote Copy configuration

In a 1-to-1 bidirectional Remote Copy configuration, each storage system in the pair functions as both the primary and secondary system, and each system contains both primary and secondary volume groups. Each system provides backup for the other, according to the direction of replication specified for each volume group.

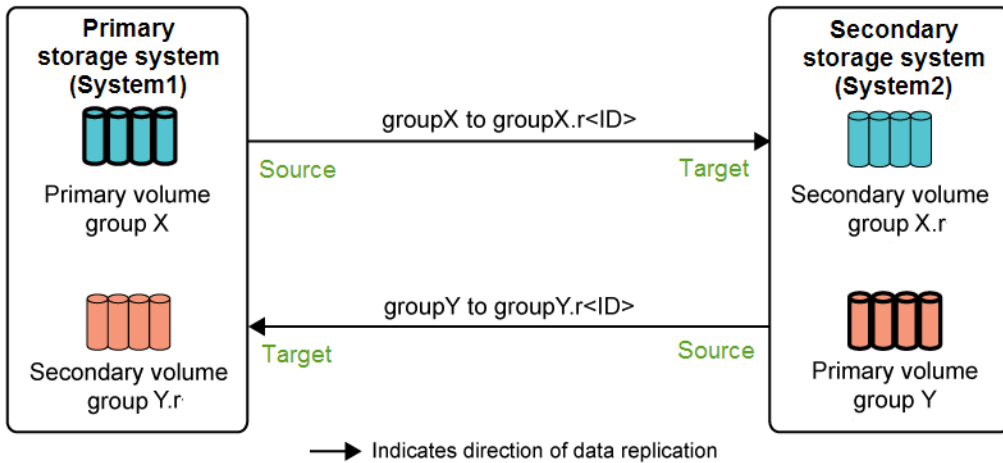


Figure 14: 1-to-1 bidirectional Remote Copy

M-to-N Remote Copy configurations

Any two storage systems in an M-to-N Remote Copy configuration are configured the same as two systems in a 1-to-1 configuration. Any single storage system in an M-to-N configuration can have a bidirectional Remote Copy relationship with up to four other storage systems. You cannot replicate the same data to more than one system in an M x N configuration.

The Remote Copy links can be RCFC or RCIP or a mixture. Only one RCIP port per node is provided; the other links may be RCFC.

Replication modes can be synchronous periodic, asynchronous periodic, or asynchronous streaming. If you want to use more than one mode, create a separate target for each mode.




Figure 15: M-to-N Remote Copy configuration example

More information

- *HPE 3PAR StoreServ 3PAR OS 3.3.1 Support Matrix*
- *HPE 3PAR StoreServ - Remote Copy Configuration Matrix*

Synchronous long-distance Remote Copy configuration

A synchronous long distance (SLD) Remote Copy configuration is composed of three systems: two systems in a synchronous Remote Copy configuration and the two synchronous systems in an asynchronous periodic relationship with the third system.

In an SLD configuration, Remote Copy volume groups from the primary system are replicated to the two separate target systems simultaneously. Data replicates to one target in synchronous mode and to the second in asynchronous periodic mode. The target systems must each have appropriate network connections to the source systems to support the replication mode being used.

Any or all storage systems in an SLD configuration can be two-node or four-node systems.

While it is not a requirement, Hewlett Packard Enterprise recommends that an SLD Remote Copy is configured:

- In synchronous mode between the primary system and the closer of the two secondary systems. The chosen systems must be within supported network latency limitations for synchronous Remote Copy.
- In asynchronous periodic mode between the primary system and the secondary system that is physically farther apart (and therefore likely to share a lesser connection).

NOTE: SLD Remote Copy links can be bidirectional between two storage systems using synchronous mode. This bidirectional synchronous relationship enables customers to configure two SLD configurations among three storage systems.

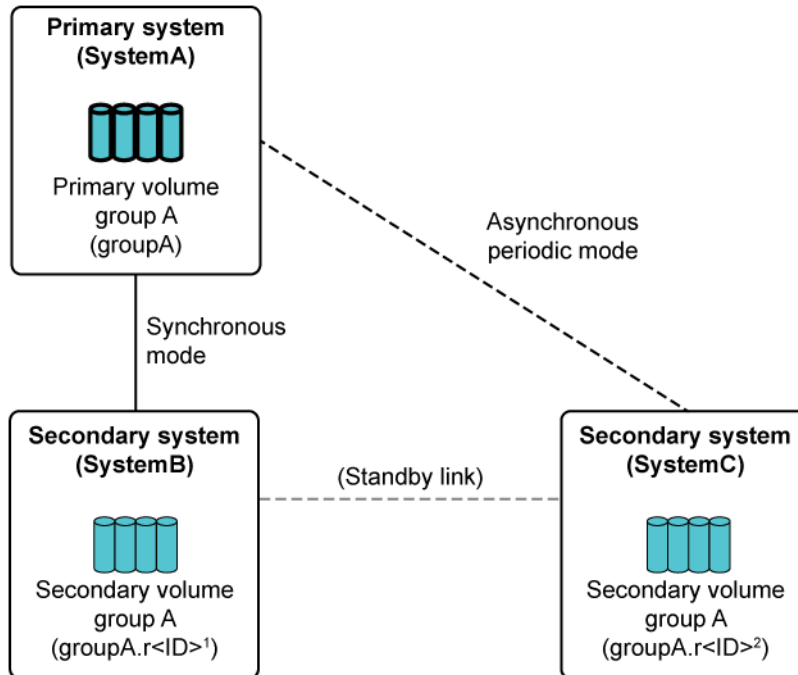


Figure 16: SLD Remote Copy

Remote Copy with HPE 3PAR Peer Persistence configuration

HPE 3PAR Peer Persistence redirects host I/O from the primary to the secondary storage system transparent to the host and without application downtime. Peer Persistence allows hosts to remain online while the servicing of their I/O requests moves from one storage system in the Peer Persistence configuration to the other.

Peer Persistence uses synchronous Remote Copy groups.

Peer Persistence offers:

- **Manual transparent failover (MTF)**

A Peer Persistence MTF or manual "switchover" operation results in:

- Failover of the host I/O for the volumes in the Remote Copy group
- Host I/O service moves from one storage system in the Peer Persistence configuration to the other
- Reversal in the direction of data replication

The switching of the servicing of the I/O from one storage system to the other is nondisruptive to the server. The switchover operation facilitates service optimization and storage system maintenance activities within a high-availability data storage solution.

- **Automatic transparent failover (ATF)**

A Peer Persistence ATF operation redirects host I/O from a failed primary system to the secondary storage system transparent to the host. An ATF executes automatically by the Peer Persistence software in response to a detected failure of the primary system for a Remote Copy group.

Peer Persistence ATF requires the configuration of a Quorum witness which provides an alternative means for the storage systems to communicate besides the Remote Copy link. This alternative communication path is necessary to ensure that a failure has occurred before an ATF can be executed.

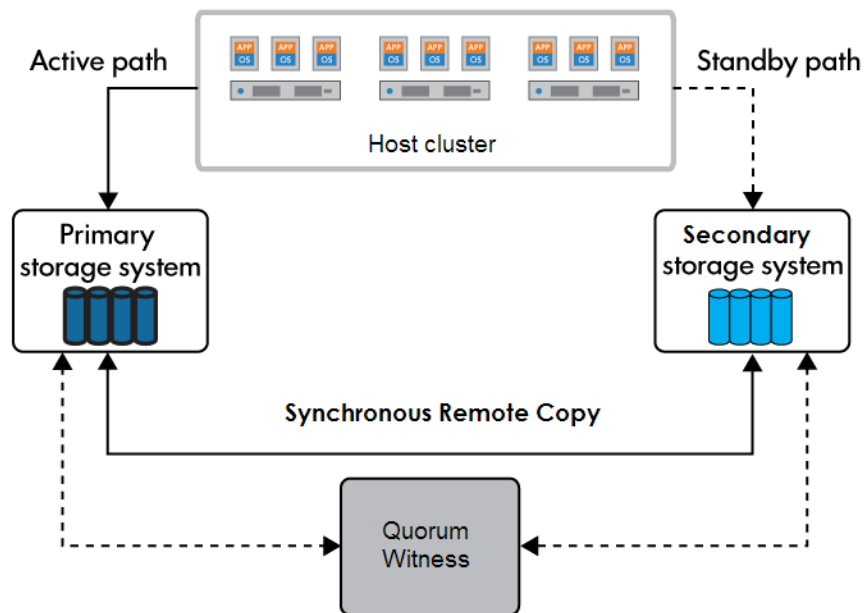


Figure 17: Peer Persistence with Quorum Witness Remote Copy

More information

HPE 3PAR Peer Persistence Host OS Support Matrix

Three Data Center Peer Persistence configuration

Three Data Center Peer Persistence (3DC-PP) enables the data from the Peer Persistence configuration to also replicate to a third storage system. The communication links and targets for the three systems are set up like an SLD topology.

- The synchronous targets are configured for Peer Persistence.
- The asynchronous targets to the third location are configured for asynchronous periodic replication.
- The Quorum Witness connections are set up the same as a standard 1-to-1 Peer Persistence configuration.

The Quorum Witness must be located on a public network and not co-located with either of the synchronous storage systems. The location of the tertiary or periodic asynchronous target system is an ideal location for the Quorum Witness, provided it is not on the same site with one of the synchronous target systems.

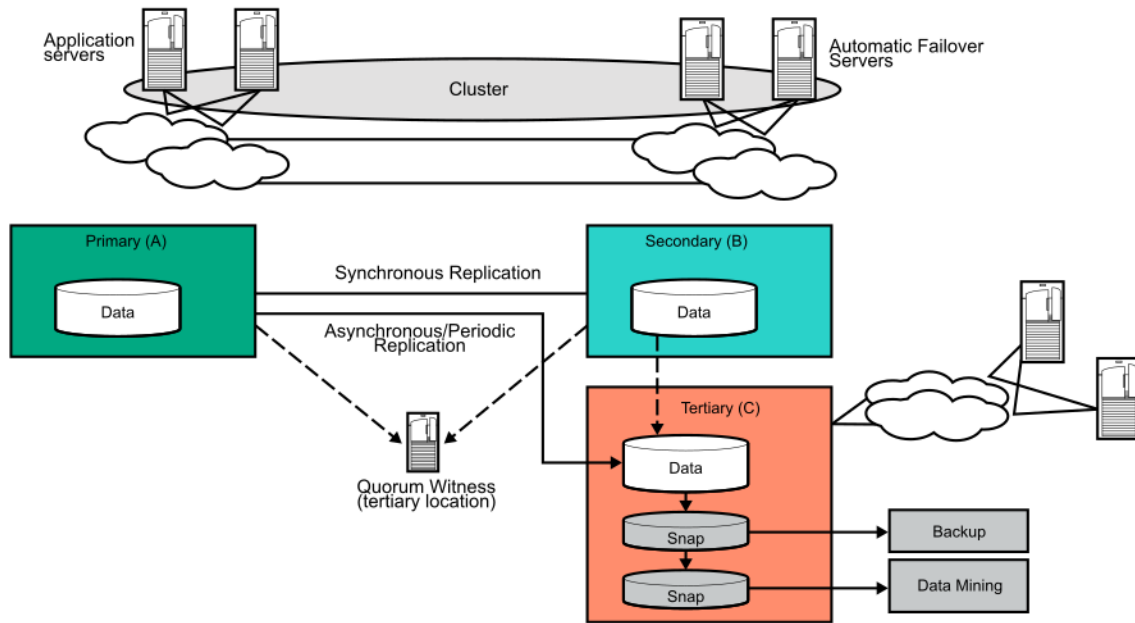


Figure 18: Three Data Center Peer Persistence (with Quorum Witness) Remote Copy

Synchronization with Remote Copy

This chapter explains the types of synchronization, the synchronization limits on virtual volumes, the use of snapshots, and how synchronization works with the different replication modes.

Types of synchronization for Remote Copy

Remote Copy performs one of two types of synchronization on Remote Copy groups:

Full synchronization

Remote Copy copies the entire primary Remote Copy group.

Initial synchronization of data is always a full synchronization when you create a Remote Copy group. Full synchronization also occurs if the administrator forces a full synchronization of the volumes in a Remote Copy group. Full synchronization occurs regardless of mode (asynchronous periodic, asynchronous streaming, or synchronous) when a group is started for the first time.

Resynchronization

Remote Copy copies only the data that has changed since the previous synchronization.

Remote Copy volumes and synchronization

To reduce the impact of synchronization on the performance of the storage systems, Remote Copy manages the number of virtual volumes that synchronize at the same time. For example, if Remote Copy is in the process of resynchronizing virtual volumes in an asynchronous periodic mode volume group, some of the first volumes complete synchronization before some of the last volumes start synchronizing.

To view resynchronizations, monitor the `SyncStatus` column of the `showrcopy` command output.

```
showrcopy
```

```
Group Information
```

Name	Target	Status	Role	Mode	Options
syncGroup1	Target_System2	Started	Primary	Sync	
LocalVV	ID	RemoteVV	ID	SyncStatus	LastSyncTime
VV1	201	VV1	92	Syncing (0%)	NA
VV2	202	VV2	93	Syncing (0%)	NA

```
showrcopy
```

```
Group Information
```

Name	Target	Status	Role	Mode	Options
syncGroup1	Target_System2	Started	Primary	Sync	
LocalVV	ID	RemoteVV	ID	SyncStatus	LastSyncTime
VV1	201	VV1	92	Syncing (60%)	NA
VV2	202	VV2	93	Syncing (60%)	NA

```
showrcopy
```

```
Group Information
```

Name	Target	Status	Role	Mode	Options
syncGroup1	Target_System2	Started	Primary	Sync	
LocalVV	ID	RemoteVV	ID	SyncStatus	LastSyncTime
VV1	201	VV1	92	Synced	NA
VV2	202	VV2	93	Synced	NA

Snapshots and coordinated snapshots within Remote Copy

A **snapshot** is a point-in-time virtual copy of a virtual volume. Under normal operating conditions, a snapshot includes only new data saved to the primary virtual volume since the last synchronization with the secondary virtual volume.

Snapshots are created to track updates to the primary volume when new data is not immediately sent to the secondary volume, like in asynchronous periodic volume groups.

To minimize the amount of data sent over the network during resynchronization, Remote Copy creates snapshots of the updates to the primary volumes. Only these snapshots are sent to secondary volumes to back up systems.

Remote Copy automatically creates snapshots when a system, or all links, go down as a result of failure or scheduled downtime.

Coordinated snapshots are snapshots that are created at the same point in time on both the primary and secondary systems. Coordinated snapshots are an important aspect of asynchronous streaming volume groups.

Synchronization with synchronous mode volume groups

Full synchronization

If the volumes in a Remote Copy group have never been synchronized, a full synchronization begins when the group is started. The length of time for the synchronization to complete depends on:

- the size of the volumes
- the amount of data the volumes contain, if they are not fully provisioned virtual volumes
- the size of the replication links
- the rate of new host writes to the virtual volumes in the group

Resynchronization

To resynchronize synchronous mode volume groups, Remote Copy copies the I/O directly from the primary base volume.

When synchronous mode volume groups are stopped, Remote Copy creates snapshots of the volume groups. When these volume groups are restarted, Remote Copy uses the snapshots to resynchronize the primary and secondary volume groups. The length of time for the resynchronization to complete depends on:

- how much data changed on the virtual volumes in the group from the time the group suspended
- the speed of the replication links
- the rate of new host writes to the virtual volumes in the group

Synchronization with asynchronous periodic mode volume groups

Initial synchronization takes a snapshot of all the virtual volumes in the Remote Copy group and replicates those snapshots to the secondary system. After the snapshot data has been replicated to the secondary system, a second set of delta snapshots is taken on the primary system. The difference between the initial snapshot and the set of delta snapshots is sent to the secondary system. The process

continues until the difference data between the snapshots can be sent within the defined delta resynchronization interval defined for the group (the "period" value of the group).

Once the initial synchronization has completed, periodic asynchronous Remote Copy enters delta resynchronization mode. At this point, Remote Copy creates delta resynchronization snapshots every "period" value and sends the delta data to the secondary system. The delta data is the changes between the previous delta resynchronization snapshot and the new delta resynchronization snapshot.

Automatic resynchronization for asynchronous periodic mode volume groups

For Remote Copy to resynchronize asynchronous periodic mode volume groups automatically, you must configure the synchronization period for the volume group. There is no default synchronization period.

When you set the synchronization period for the primary volume group, the network link used for replication must be large enough to ensure that all delta data can be transmitted within the specified synchronization interval.

The resynchronization schedule starts its countdown from the current resynchronization start time. Therefore, if a resynchronization is delayed, the next resynchronization starts after the interval specified as the synchronization period.

For example, if you set the synchronization period to 10 minutes and one synchronization begins at 1 pm (and completes within the 10 minute synchronization period), the next synchronization begins at 1:10 pm, the next at 1:20 pm, and so on. If you set the synchronization period to 10 minutes, and one resynchronization begins at 1 pm but takes 12 minutes to complete, the next resynchronization begins around 1:12 pm. If that resynchronization takes only 2 minutes to complete (which is within the 10 minute interval), the next resynchronization begins at 1:22, 10 minutes after the previous resynchronization.

You can adjust the synchronization period any time after you create the volume group, even after you start Remote Copy operations for the group.

Synchronization with asynchronous streaming mode

Full synchronization

If the volumes in a Remote Copy group have never been synchronized, a full synchronization begins when the group is started.

While the primary and secondary volume groups are synchronizing again, the primary system accepts I/O and sends it to the secondary system. The ordering between host I/Os and initial synchronization I/Os is preserved.

Resynchronization

Periodically (hourly by default), resynchronization points (coordinated snapshots) are created on the primary and secondary Remote Copy groups. After asynchronous streaming mode volume groups are stopped, the last coordinated snapshot is used to resynchronize the primary and secondary volume groups when the groups are restarted.

Snapshot frequency for asynchronous streaming mode volume groups

For Remote Copy to resynchronize asynchronous streaming mode volume groups quickly, you must configure the synchronization snapshot frequency for the volume group.

You can set the synchronization snapshot frequency any time after you create the volume group, even after you start Remote Copy operations for the group.

When you set the synchronization snapshot frequency, you specify the frequency using the `setrcopygroup snap_freq` command.

Resynchronization order for asynchronous streaming mode volume groups

For asynchronous streaming Remote Copy groups, you can specify the order the groups will be stopped and restarted after the replication resources are restored. The order is determined by the "period" value assigned to the Remote Copy group. The "period" value also specifies how long the asynchronous streaming group remains stopped before Remote Copy tries to restart it. If resources are determined to be insufficient to restart the group, it will wait another "period" value before trying to restart again.

Groups with larger period values will be stopped first and restarted last. Setting the value allows you to prioritize your asynchronous streaming Remote Copy groups. For example, assign a low value to your critical groups to bring them back online before less critical groups. HPE recommends that you do not set the same value to more than one group.

Setting a period value of zero (0) will suspend the group first and will not automatically restart the group. The group must be manually restarted if Remote Copy suspends the group.

System access, requirements, and best practices for Remote Copy

This chapter explains system access, prerequisites, and best practices in using Remote Copy.

System access to set up Remote Copy

To set up Remote Copy, you must have access to one of the following tools:

- **HPE 3PAR StoreServ Management Console SSMC** where you can use the Remote Copy Setup wizard.
- **HPE 3PAR CLI commands** as described in the *HPE 3PAR Remote Copy Software User Guide*.

NOTE: As of SSMC 3.2, the SSMC supports all configurations and modes listed in the support matrix for HPE 3PAR OS 3.3.1. However, SSMC is recommended for basic 1-to-1 configurations. CLI commands are recommended for more complex Remote Copy configurations, such as M-to-N and SLD.

Remote Copy requirements and best practices

Before you set up your Remote Copy system, you must meet the requirements and should understand the best practices.

Prerequisites for setting up Remote Copy

Ensure the following requirements are met before setting up your Remote Copy system:

- **At least two HPE 3PAR StoreServ Storage systems.** You have access to a minimum of two HPE 3PAR StoreServ Storage systems which are supported for Remote Copy.
See the *HPE 3PAR Support Matrix* and the *HPE 3PAR Operating System Upgrade Planning Guide* based on your HPE 3PAR OS version.
- **HPE 3PAR All-inclusive Multi-System Software License** for all storage systems participating in Remote Copy replication.
See the *HPE 3PAR StoreServ Storage Concepts Guide*.
- **IP or Fibre Channel networks.** The physical connections between all storage systems used with Remote Copy are through an IP-capable network or a Fibre Channel (FC) SAN network.
See the *SAN design Reference Guide* at HPE Support: www.hpe.com/support/hpesc.
- **Firewall allows traffic.** Your firewall allows traffic over the designated Remote Copy over IP (RCIP) only ports, TCP 5785 and 5001.
- **Unique IP addresses.** All RCIP ports have unique IP addresses. For example: 10.100.32.0, 10.100.32.1, 09.101.90.45.
- **SSMC or the HPE 3PAR CLI.** You have access to SSMC or the HPE 3PAR CLI.
This guide covers CLI commands. For SSMC, see the *HPE 3PAR StoreServ Management Console 3.1 User Guide*.

Software restrictions for Remote Copy

- Remote Copy does not support self-mirroring configurations. It cannot use a storage system to replicate its own primary volumes to itself.
- Remote Copy does not support multi-hop configurations. It cannot replicate a primary volume group to a secondary system and then replicate the volume group again from the secondary system to a third storage system.

Best practices for Remote Copy

Consider the following best practices before setting up your Remote Copy system.

- **Distance between the systems**

- Storage systems in the same room can be connected through network adapter (NIC) switches, or through FC networks. Longer distances require other topologies.
- For disaster recovery purposes, the distance between storage systems depends on the disaster from which you need to recover (for example, a local power outage versus a widespread natural disaster).
- In synchronous replication mode, the latency of Remote Copy writes increases with distance.

- **Bandwidth**

RCIP configurations can fully use all 1 Gbps or 10Gbps links during the synchronization process. If you use networks with lower throughput, the throughput between systems running Remote Copy software might be capped by the network speed.

- **Connection requirements**

Network connections (or transport methods) can be IP or FC.

- **Data availability and speed**

- To maintain data availability, you must use at least two links to connect storage systems.
- Remote Copy requires at least two separate network connections between storage systems to ensure availability in case one of the networks fails. For optimal availability, connect the two HPE 3PAR StoreServ Storage systems to separate network equipment.
- When there are multiple physical links between the storage systems, Remote Copy uses all the available links that are configured for the same Remote Copy target to transmit data in parallel.
- To limit the potential for data loss, Hewlett Packard Enterprise recommends that synchronous mode replication be used whenever the additional write latency will not exceed the maximum write latency tolerable by the application whose data is being replicated.

Best Practice for Remote Copy volumes

Consider the following best practices for virtual volumes used in Remote Copy.

- Source and target virtual volumes must be the same size.
- Virtual volumes must have a user CPG and a snapshot CPG.
- A snapshot CPG for virtual volumes being replicated must have the same tier of storage as the source; not a lower tier of storage.

- To provide the best granularity of Remote Copy failover, allow a failover of only one host or application. Therefore, add all virtual volumes belonging to the same application to the same Remote Copy group. Do not add virtual volumes that have no host or application (logical link) to the same Remote Copy group.

Virtual volumes that have logical links include:

- Virtual volumes used by the same application (data and log volumes, for example)
 - Virtual volumes used by a Logical Volume Manager (LVM) volume group
 - Virtual volumes that contain virtual disks of the same VM
- Consider performance before mixing different types of disks/RAID levels on the source and target systems. You can only write as fast as your slowest drive.
 - When using Remote Copy in combination with virtual domains, the source and target virtual volumes must be in domains of the same name on the source and target systems. However, they do not need to be of the same type (type of provisioning, RAID, or disk).

Documentation, support, and other resources

This part provides a list of other documentation referenced in this guide and their location. Support information is also provided.

Supporting documentation for Remote Copy

The following documents are referenced in this guide and provide additional information for your Remote Copy solution.

Document	Location
<i>HPE 3PAR StoreServ 3PAR OS 3.3.1 Support Matrix</i>	Single Point of Connectivity Knowledge (SPOCK) http://www.hpe.com/storage/spock
<i>HPE 3PAR Remote Copy Configuration Matrix</i>	Single Point of Connectivity Knowledge (SPOCK) http://www.hpe.com/storage/spock
<i>Disaster-tolerant solutions with HPE 3PAR Remote Copy white paper</i>	Hewlett Packard Enterprise Information Library http://www.hpe.com/info/storage/docs
<i>SAN design Reference Guide</i>	HPE Support Center http://www.hpe.com/support/hpesc
<i>HPE 3PAR Command Line Interface Administrator Guide</i>	Hewlett Packard Enterprise Information Library http://www.hpe.com/info/storage/docs
<i>HPE 3PAR Command Line Interface Reference</i>	Hewlett Packard Enterprise Information Library http://www.hpe.com/info/storage/docs
<i>HPE 3PAR StoreServ Management Console User Guide</i>	Hewlett Packard Enterprise Information Library http://www.hpe.com/info/storage/docs
<i>HPE 3PAR Operating System Upgrade Planning Guide</i>	Hewlett Packard Enterprise Information Library http://www.hpe.com/info/storage/docs
<i>HPE 3PAR StoreServ Storage Concepts Guide</i>	Hewlett Packard Enterprise Information Library http://www.hpe.com/info/storage/docs
<i>HPE 3PAR Peer Persistence Host OS Support Matrix</i>	Single Point of Connectivity Knowledge (SPOCK) http://www.hpe.com/storage/spock

Support and other resources

Accessing Hewlett Packard Enterprise Support

- For live assistance, go to the Contact Hewlett Packard Enterprise Worldwide website:
<http://www.hpe.com/assistance>
- To access documentation and support services, go to the Hewlett Packard Enterprise Support Center website:
<http://www.hpe.com/support/hpesc>

Information to collect

- Technical support registration number (if applicable)
- Product name, model or version, and serial number
- Operating system name and version
- Firmware version
- Error messages
- Product-specific reports and logs
- Add-on products or components
- Third-party products or components

Accessing updates

- Some software products provide a mechanism for accessing software updates through the product interface. Review your product documentation to identify the recommended software update method.

- To download product updates:

Hewlett Packard Enterprise Support Center

www.hpe.com/support/hpesc

Hewlett Packard Enterprise Support Center: Software downloads

www.hpe.com/support/downloads

Software Depot

www.hpe.com/support/softwaredepot

- To subscribe to eNewsletters and alerts:

www.hpe.com/support/e-updates

- To view and update your entitlements, and to link your contracts and warranties with your profile, go to the Hewlett Packard Enterprise Support Center **More Information on Access to Support Materials** page:

www.hpe.com/support/AccessToSupportMaterials

! **IMPORTANT:** Access to some updates might require product entitlement when accessed through the Hewlett Packard Enterprise Support Center. You must have an HPE Passport set up with relevant entitlements.

Customer self repair

Hewlett Packard Enterprise customer self repair (CSR) programs allow you to repair your product. If a CSR part needs to be replaced, it will be shipped directly to you so that you can install it at your convenience. Some parts do not qualify for CSR. Your Hewlett Packard Enterprise authorized service provider will determine whether a repair can be accomplished by CSR.

For more information about CSR, contact your local service provider or go to the CSR website:

<http://www.hpe.com/support/selfrepair>

Remote support

Remote support is available with supported devices as part of your warranty or contractual support agreement. It provides intelligent event diagnosis, and automatic, secure submission of hardware event notifications to Hewlett Packard Enterprise, which will initiate a fast and accurate resolution based on your product's service level. Hewlett Packard Enterprise strongly recommends that you register your device for remote support.

If your product includes additional remote support details, use search to locate that information.

Remote support and Proactive Care information

HPE Get Connected

www.hpe.com/services/getconnected

HPE Proactive Care services

www.hpe.com/services/proactivecare

HPE Proactive Care service: Supported products list

www.hpe.com/services/proactivecaresupportedproducts

HPE Proactive Care advanced service: Supported products list

www.hpe.com/services/proactivecareadvancedsupportedproducts

Proactive Care customer information

Proactive Care central

www.hpe.com/services/proactivecarecentral

Proactive Care service activation

www.hpe.com/services/proactivecarecentralgetstarted

Warranty information

To view the warranty information for your product, see the links provided below:

HPE ProLiant and IA-32 Servers and Options

www.hpe.com/support/ProLiantServers-Warranties

HPE Enterprise and Cloudline Servers

www.hpe.com/support/EnterpriseServers-Warranties

HPE Storage Products

www.hpe.com/support/Storage-Warranties

HPE Networking Products

www.hpe.com/support/Networking-Warranties

Regulatory information

To view the regulatory information for your product, view the *Safety and Compliance Information for Server, Storage, Power, Networking, and Rack Products*, available at the Hewlett Packard Enterprise Support Center:

www.hpe.com/support/Safety-Compliance-EnterpriseProducts

Additional regulatory information

Hewlett Packard Enterprise is committed to providing our customers with information about the chemical substances in our products as needed to comply with legal requirements such as REACH (Regulation EC No 1907/2006 of the European Parliament and the Council). A chemical information report for this product can be found at:

www.hpe.com/info/reach

For Hewlett Packard Enterprise product environmental and safety information and compliance data, including RoHS and REACH, see:

www.hpe.com/info/ecodata

For Hewlett Packard Enterprise environmental information, including company programs, product recycling, and energy efficiency, see:

www.hpe.com/info/environment

Documentation feedback

Hewlett Packard Enterprise is committed to providing documentation that meets your needs. To help us improve the documentation, send any errors, suggestions, or comments to Documentation Feedback (docsfeedback@hpe.com). When submitting your feedback, include the document title, part number, edition, and publication date located on the front cover of the document. For online help content, include the product name, product version, help edition, and publication date located on the legal notices page.