

Cluster Server Agent for Hitachi/HP XP 3 Data Center Replication Configuration Guide

HP-UX

5.0

Cluster Server Agent for Hitachi/HP XP 3 Data Center Replication Configuration Guide

The software described in this book is furnished under a license agreement and may be used only in accordance with the terms of the agreement.

Agent Version: 5.0

Document version: 5.0 Rev 0

Legal Notice

Copyright © 2015 Symantec Corporation. All rights reserved.

Symantec, the Symantec Logo, the Checkmark Logo, Veritas, the Veritas Logo, InfoScale, CommandCentral, NetBackup, Enterprise Vault, and LiveUpdate are trademarks or registered trademarks of Symantec Corporation or its affiliates in the U.S. and other countries. Other names may be trademarks of their respective owners.

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 Symantec Corporation 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. SYMANTEC CORPORATION 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, "Rights in Commercial Computer Software or Commercial Computer Software Documentation", as applicable, and any successor regulations. 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.

Symantec Corporation
350 Ellis Street
Mountain View, CA 94043

<http://www.symantec.com>

Technical Support

Symantec Technical Support maintains support centers globally. Technical Support's primary role is to respond to specific queries about product features and functionality. The Technical Support group also creates content for our online Knowledge Base. The Technical Support group works collaboratively with the other functional areas within Symantec to answer your questions in a timely fashion. For example, the Technical Support group works with Product Engineering and Symantec Security Response to provide alerting services and virus definition updates.

Symantec's support offerings include the following:

- A range of support options that give you the flexibility to select the right amount of service for any size organization
- Telephone and/or Web-based support that provides rapid response and up-to-the-minute information
- Upgrade assurance that delivers software upgrades
- Global support purchased on a regional business hours or 24 hours a day, 7 days a week basis
- Premium service offerings that include Account Management Services

For information about Symantec's support offerings, you can visit our website at the following URL:

www.symantec.com/business/support/index.jsp

All support services will be delivered in accordance with your support agreement and the then-current enterprise technical support policy.

Contacting Technical Support

Customers with a current support agreement may access Technical Support information at the following URL:

www.symantec.com/business/support/contact_techsupp_static.jsp

Before contacting Technical Support, make sure you have satisfied the system requirements that are listed in your product documentation. Also, you should be at the computer on which the problem occurred, in case it is necessary to replicate the problem.

When you contact Technical Support, please have the following information available:

- Product release level
- Hardware information

- Available memory, disk space, and NIC information
- Operating system
- Version and patch level
- Network topology
- Router, gateway, and IP address information
- Problem description:
 - Error messages and log files
 - Troubleshooting that was performed before contacting Symantec
 - Recent software configuration changes and network changes

Licensing and registration

If your Symantec product requires registration or a license key, access our technical support Web page at the following URL:

www.symantec.com/business/support/

Customer service

Customer service information is available at the following URL:

www.symantec.com/business/support/

Customer Service is available to assist with non-technical questions, such as the following types of issues:

- Questions regarding product licensing or serialization
- Product registration updates, such as address or name changes
- General product information (features, language availability, local dealers)
- Latest information about product updates and upgrades
- Information about upgrade assurance and support contracts
- Information about the Symantec Buying Programs
- Advice about Symantec's technical support options
- Nontechnical presales questions
- Issues that are related to CD-ROMs or manuals

Support agreement resources

If you want to contact us regarding an existing support agreement, please contact the support agreement administration team for your region as follows:

Asia-Pacific and Japan customercare_apj@symantec.com

Europe, Middle-East, and Africa semea@symantec.com

North America and Latin America [supportolutions@symantec.com](mailto:supportsolutions@symantec.com)

Documentation

Product guides are available on the media in PDF format. Make sure that you are using the current version of the documentation. The document version appears on page 2 of each guide. The latest product documentation is available on the Symantec website.

<https://sort.symantec.com/documents>

Your feedback on product documentation is important to us. Send suggestions for improvements and reports on errors or omissions. Include the title and document version (located on the second page), and chapter and section titles of the text on which you are reporting. Send feedback to:

doc_feedback@symantec.com

For information regarding the latest HOWTO articles, documentation updates, or to ask a question regarding product documentation, visit the Storage and Clustering Documentation forum on Symantec Connect.

<https://www-secure.symantec.com/connect/storage-management/forums/storage-and-clustering-documentation>

About Symantec Connect

Symantec Connect is the peer-to-peer technical community site for Symantec's enterprise customers. Participants can connect and share information with other product users, including creating forum posts, articles, videos, downloads, blogs and suggesting ideas, as well as interact with Symantec product teams and Technical Support. Content is rated by the community, and members receive reward points for their contributions.

<http://www.symantec.com/connect/storage-management>

Contents

Technical Support	3	
Chapter 1	Introducing the agent for Hitachi/HP XP 3 Data Center Replication	8
	About the agent for Hitachi/HP XP 3 Data Center Replication	8
	Configuration Type 1	9
	Configuration Type 2	10
	Supported software	13
	Supported hardware for Hitachi/HP XP 3 Data Center Replication	13
	Typical Hitachi/HP XP 3 Data Center Replication setup in a VCS cluster	13
	Configuration Type 1	14
	Configuration Type 2	15
	Agent functions	22
	About the agent's online function	25
Chapter 2	Installing and removing the agent for Hitachi/HP XP 3 Data Center Replication	28
	Before you install the agent for Hitachi/HP XP 3 Data Center Replication	28
	Installing the agent for Hitachi/HP XP 3 Data Center Replication	28
	Removing the agent for Hitachi/HP XP 3 Data Center Replication	29
Chapter 3	Configuring the agent for Hitachi/HP XP 3 Data Center Replication	30
	Configuration concepts for the Hitachi/HP XP 3 Data Center Replication agent	30
	Resource type definition for the Hitachi/HP XP 3DC agent	30
	Attribute definitions for the Hitachi/HP XP 3DC agent	31
	Sample configuration for the Hitachi/HP XP 3DC agent	37
	Before you configure the agent for Hitachi/HP XP 3 Data Center Replication	39
	About cluster heartbeats	39

	Configuring the agent for Hitachi/HP XP 3 Data Center	
	Replication	39
	Configuring the agent to compute RPO	40
Chapter 4	Testing VCS disaster recovery support with Hitachi/HP XP 3 Data Center Replication	42
	Testing service group migration from PRIMARY to SYNC_TARGET	42
	Failure scenarios at PRIMARY	43
	Production cluster failure	48
	Performing failback from ASYNC_TARGET	49
Index		50

Introducing the agent for Hitachi/HP XP 3 Data Center Replication

This chapter includes the following topics:

- [About the agent for Hitachi/HP XP 3 Data Center Replication](#)
- [Supported software](#)
- [Supported hardware for Hitachi/HP XP 3 Data Center Replication](#)
- [Typical Hitachi/HP XP 3 Data Center Replication setup in a VCS cluster](#)
- [Agent functions](#)

About the agent for Hitachi/HP XP 3 Data Center Replication

The Cluster Server agent for Hitachi/HP XP 3 Data Center Replication provides support for application failover and recovery. The agent provides this support in environments that use Hitachi/HP XP 3DC replication to replicate data between Hitachi/HP-XP arrays.

The agent monitors and manages the state of replicated Hitachi/HP XP 3 Data Center Replication devices that are attached to VCS nodes. The agent ensures that the system that has the Hitachi/HP XP 3 Data Center Replication resource online also has safe and exclusive access to the configured devices.

For synchronous replication, the agent supports the DATA and NEVER replication fence levels.

The agent uses the following terminology:

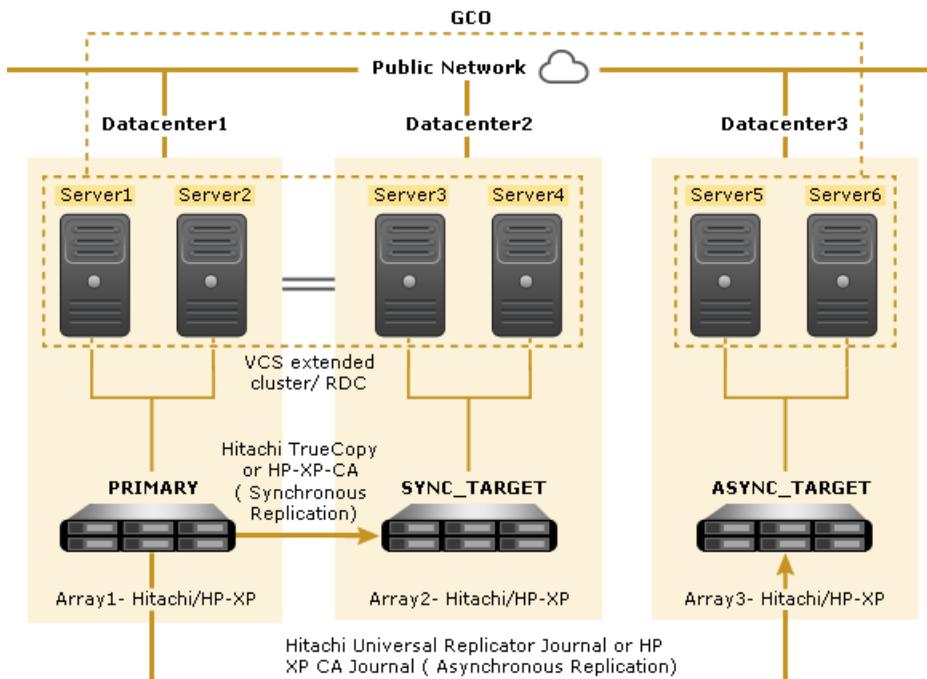
- **PRIMARY:** The datacenter that serves as the production environment. This is typically a part of a VCS extended cluster or a VCS replicated data cluster (RDC).
- **SYNC_TARGET:** The datacenter to which the production data is replicated synchronously. This is typically a part of a VCS extended cluster or a VCS RDC.
- **ASYNC_TARGET:** The datacenter to which the production data is replicated asynchronously.

The Cluster Server agent for Hitachi/HP XP 3 Data Center Replication supports the 3 Data Center (3DC) Multi-target implementation in two types of configurations:

Configuration Type 1

Figure 1-1 shows how the agent for Hitachi/HP XP 3 Data Center Replication supports the Multi-target mode of operation.

Figure 1-1 Configuration Type 1



In a typical 3DC Multi-target implementation of Configuration Type 1, data from the source volume is replicated synchronously to a nearby secondary site, and simultaneously, data from that same source volume is replicated asynchronously to a geographically remote site. This mode of operation protects production data against both local and wide-area disasters. There is no third link created between the nearby secondary and the geographically distant remote site.

In this type of configuration, after the failover to the nearby secondary is complete, the configuration changes to a 3DC Cascaded mode of replication. The agent supports failover in this 3DC Cascaded implementation as well.

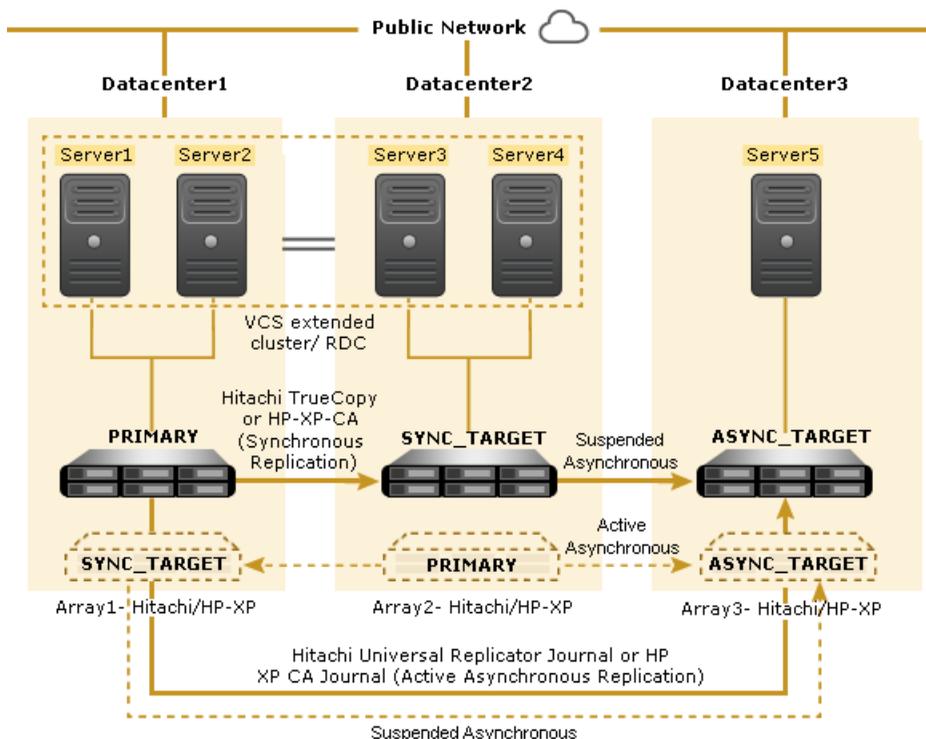
In this document, this type of configuration is referred to as Configuration Type 1.

Configuration Type 2

Configuration Type 2 is a 3DC Multi-target implementation where, in addition to data from the source volume being replicated synchronously to the nearby secondary site and asynchronously to the geographically distant remote site, there is also a suspended asynchronous replication link between the nearby secondary site and the geographically distant remote site.

[Figure 1-2](#) shows how the agent for Hitachi/HP XP 3 Data Center Replication supports this Multi-target mode of operation.

Figure 1-2 Configuration Type 2



In this configuration, in order to properly migrate an application service across the two sites in a consistent manner, it is required to associate a different journal group for each service group to be configured. However, increasing the number of journal groups is likely to impact the replication performance. In keeping with HP XP Continuous Access best practices, in order to configure more than 16/32 application services, it is required to have more than one storage.

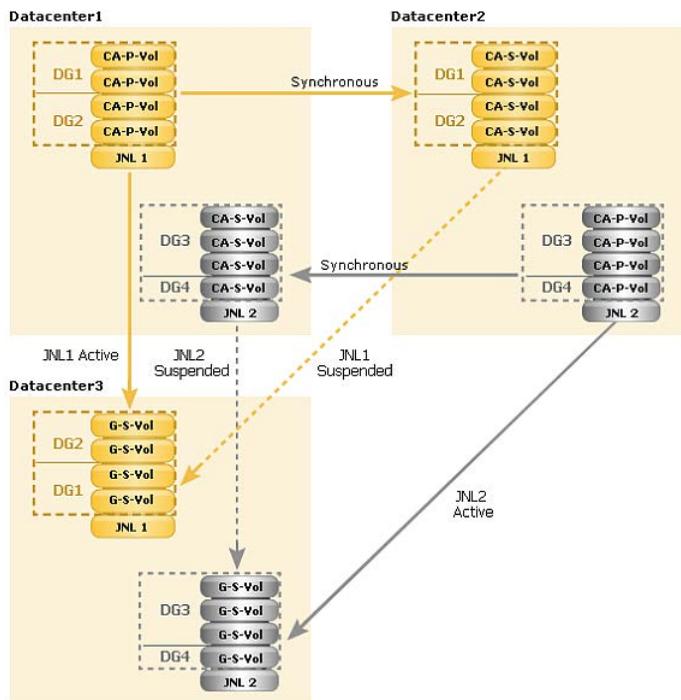
To overcome this limit, this configuration type dictates using a limited number of journal groups, typically 2 journal groups — one which includes device groups replicating from one direction and the other replicating device groups in the reverse direction.

In [Figure 1-2](#), the solid arrows represent one replication direction of the device groups and the dotted arrows represent the reverse replication direction of the other device groups.

The agent facilitates the planned migration of service groups which have replication configured in this manner.

Figure 1-3 depicts an example of an implementation where replication is configured in this manner.

Figure 1-3



Consider the example of an implementation where two different journal groups—JNL 1 and JNL 2—are configured.

- Journal group JNL 1 consists of device groups DG1 and DG2, and originates from datacenter1 to datacenter2 (synchronous link) and datacenter 1 > datacenter 3 (asynchronous PAIR link) and datacenter 2 > datacenter 3 (asynchronous PSUS link)

Where, datacenter 1 is PRIMARY, datacenter 2 is SYNC_TARGET, and datacenter 3 is ASYNC_TARGET
- Journal group JNL 2 consists of device groups DG3 and DG4, and originates from datacenter 2 > datacenter 1 (synchronous link) and datacenter 2 > datacenter 3(asynchronous PAIR link) and datacenter 1>datacenter 3 (asynchronous PSUS link)

Where, datacenter 2 is PRIMARY, datacenter 1 is SYNC_TARGET, and datacenter 3 is ASYNC_TARGET .

Each device group is then mapped to a different Volume Manager disk group and consequently to a different service group.

Supported software

For information on the software versions that the agent for Hitachi/HP XP 3 Data Center Replication supports, see the Symantec Operations Readiness Tools (SORT) site: <https://sort.symantec.com/agents>.

Supported hardware for Hitachi/HP XP 3 Data Center Replication

The agent for Hitachi/HP XP 3 Data Center Replication supports all Hitachi/HP XP arrays with microcode levels that support the 3DC configuration.

Typical Hitachi/HP XP 3 Data Center Replication setup in a VCS cluster

This topic provides additional details about the Hitachi/HP XP 3 DC Multi-target configuration types that the agent supports. Further, this topic explains how failover works in those environments.

Clustering in a Hitachi/HP XP 3DC environment typically consists of the following hardware infrastructure:

- Datacenter 1 (PRIMARY), that is, the primary array (Array1) has one or more P-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to the Hitachi/HP-XP array that contains the TrueCopy/HP XP CA/HUR Journal P-VOL devices.
- Datacenter 2 (SYNC_TARGET), that is, the secondary array (Array2) has one or more S-VOL hosts. A Fibre Channel or SCSI directly attaches these hosts to a Hitachi/HP-XP array that contains the TrueCopy/HP XP CA (Synchronous) S-VOL devices. The S-VOL devices are paired with the P-VOL devices in the synchronous device group, in Array1. The S-VOL hosts and arrays must be at a distance that is appropriate for synchronous replication.
- Datacenter 3 (ASYNC_TARGET), that is, the remote secondary array (Array3) has one or more S-VOL hosts that belong to the asynchronous device group. A Fibre Channel or SCSI directly attaches these hosts to a Hitachi/HP XP array. These S-VOL devices are paired with P-VOL devices at Datacenter 1 (PRIMARY) using HUR Journal or HP XP CA Journal replication. These S-VOL hosts and

arrays must be at a significant distance to survive a disaster at both Datacenter 1 and Datacenter 2.

- Network heartbeating between the two datacenters to determine their health; this network heartbeating could be LLT or TCP/IP.
- In a replicated data cluster environment, comprising Datacenter 1 (PRIMARY) and Datacenter 2 (SYNC_TARGET), all hosts are part of the same cluster. You must connect them with the dual and dedicated networks that support LLT.

Configuration Type 1

Figure 1-1 displays a typical cluster setup in a Hitachi/HP XP 3DC environment, where:

- PRIMARY and SYNC_TARGET form a VCS extended cluster or VCS RDC.
- VCS Global Cluster Option (GCO) is configured between the production cluster (includes PRIMARY and SYNC_TARGET) and the DR cluster (includes ASYNC_TARGET).
- In a global cluster environment at the Disaster-Recovery (DR) site, that is, Datacenter 3 (ASYNC_TARGET), you must attach all hosts in a cluster to the same Hitachi/HP-XP array.

Note: The agent supports the 3DC Multi-target configuration of Configuration Type 1, when there is no third link configured between SYNC_TARGET and ASYNC_TARGET, for example, a link to leverage the 3DC Delta Resync feature.

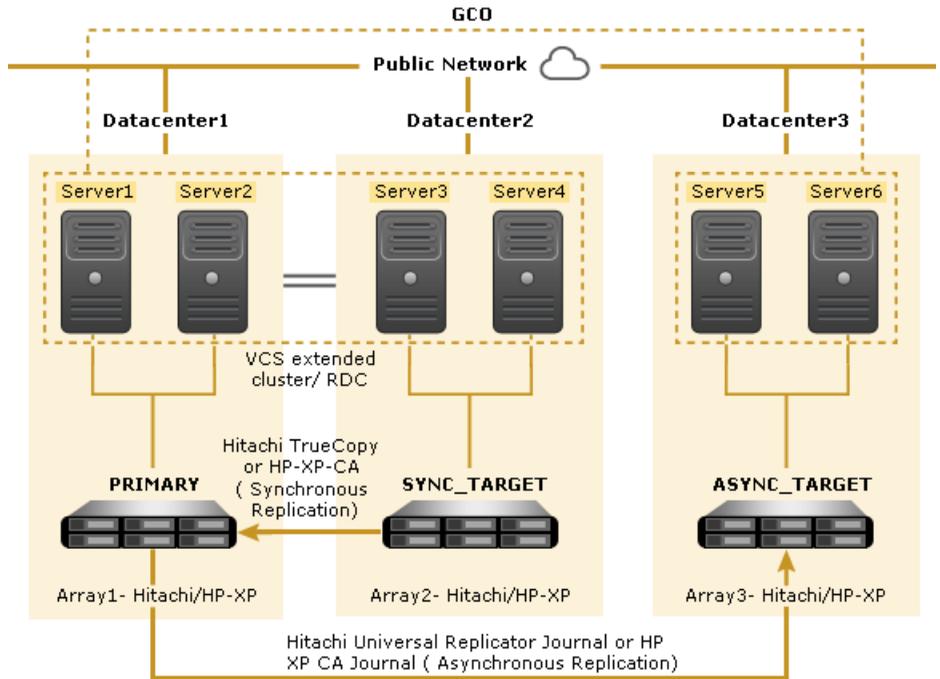
Configuration Type 1: Failover of a service

In a 3DC Multi-target mode of operation in Configuration Type 1, production data from the PRIMARY is replicated synchronously to SYNC_TARGET, which is part of the replicated data cluster. Simultaneously, production data from the PRIMARY is replicated asynchronously to the remote ASYNC_TARGET, which is part of the global cluster.

In case of a failure at the primary site or a planned migration, the agent initiates a failover to SYNC_TARGET. If there is a disaster at the production site, the application is failed over to the DR site, that is, ASYNC_TARGET. In this case, the agent initiates a takeover procedure on ASYNC_TARGET. When the production cluster is up again, the application needs to be failed back to PRIMARY. The agent takes the required actions to fail back from ASYNC_TARGET to PRIMARY.

Figure 1-4 shows how the configuration looks when failover is successfully accomplished in case of a planned migration from the production cluster.

Figure 1-4 3DC Cascaded mode of operation



If failover to SYNC_TARGET is successful, the roles are reversed – SYNC_TARGET becomes the new PRIMARY and the original PRIMARY becomes the SYNC_TARGET. In such a case, the agent employs the 3DC Cascaded mode of operation. Data is synchronously replicated from the new PRIMARY (original SYNC_TARGET) to the new SYNC_TARGET (original PRIMARY). This new SYNC_TARGET forwards the data to the remote ASYNC_TARGET, using asynchronous replication.

Configuration Type 2

Figure 1-2 displays a typical cluster setup in a Hitachi/HP XP 3DC environment, where:

- PRIMARY and SYNC_TARGET form a VCS extended cluster or VCS RDC
- A suspended asynchronous replication link exists between SYNC_TARGET and ASYNC_TARGET.

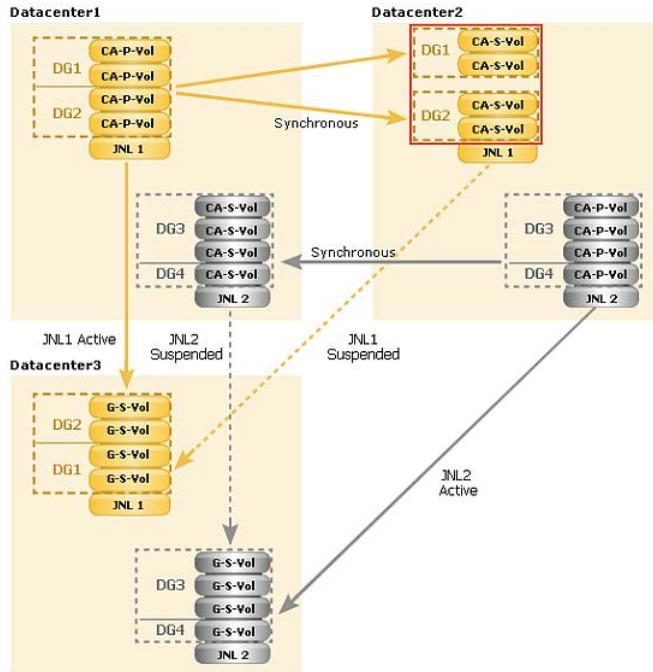
Note: The agent does not support fail over or switch of services from PRIMARY and SYNC_TARGET to ASYNC_TARGET.

Configuration Type 2: Migrating device groups across sites

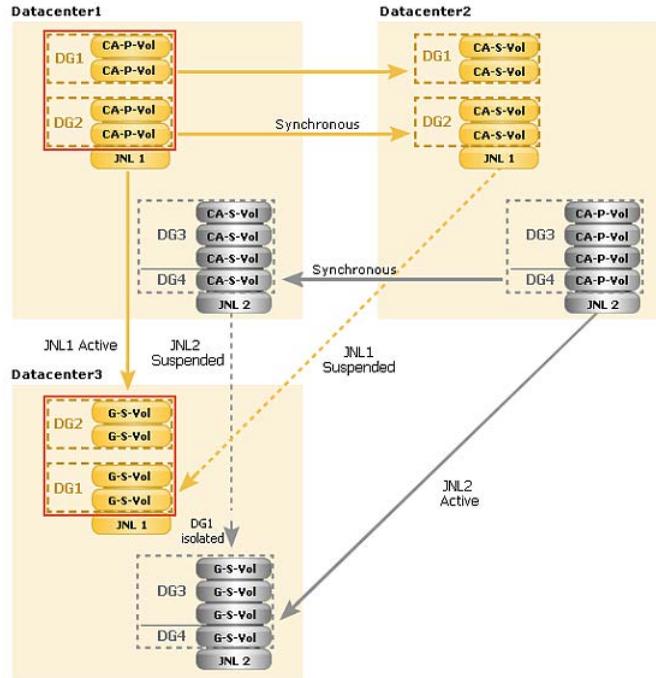
The following example shows how a device group (DG1) is migrated from a journal group (JNL 1) to another journal group (JNL 2) as a result of a planned migration.

Migrating device groups

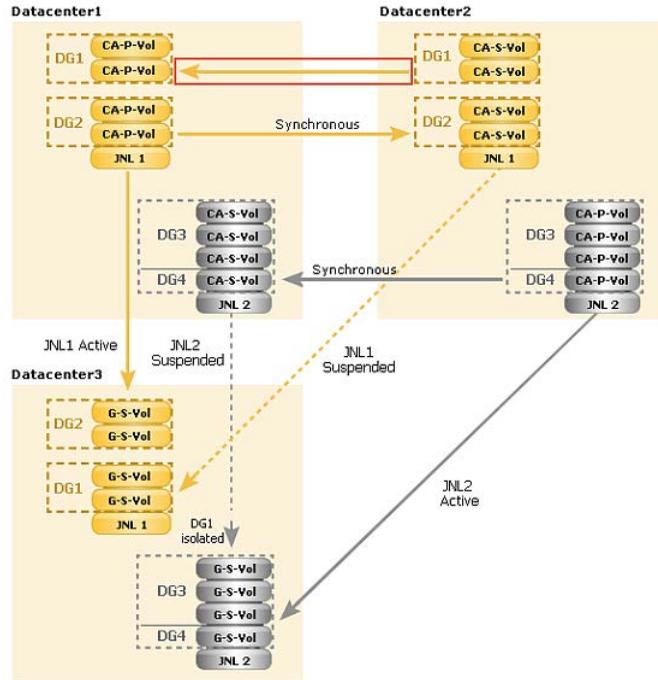
- 1 On SYNC_TARGET, the agent executes the `pairsplit -s` command to delete the suspended JNL 1 replication link for DG1.



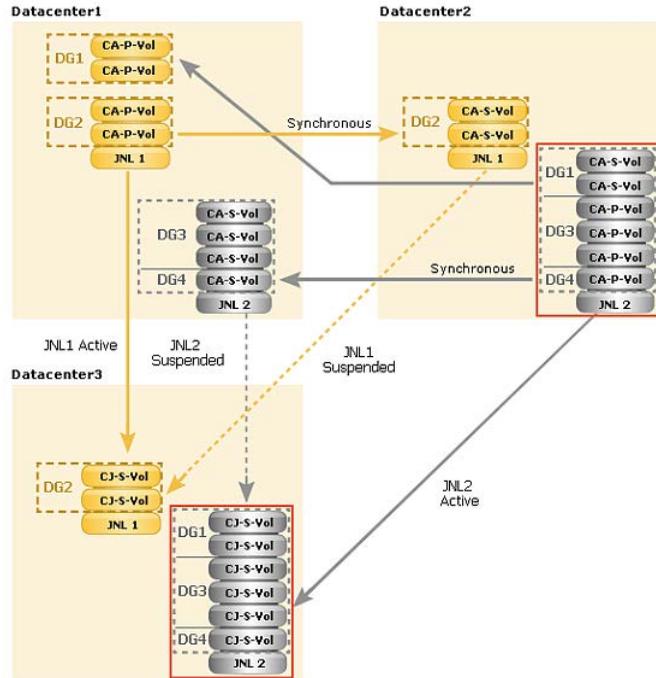
- 2 On PRIMARY, the agent executes the `pairsplit -s` command to delete the active JNL 1 replication link for DG1.



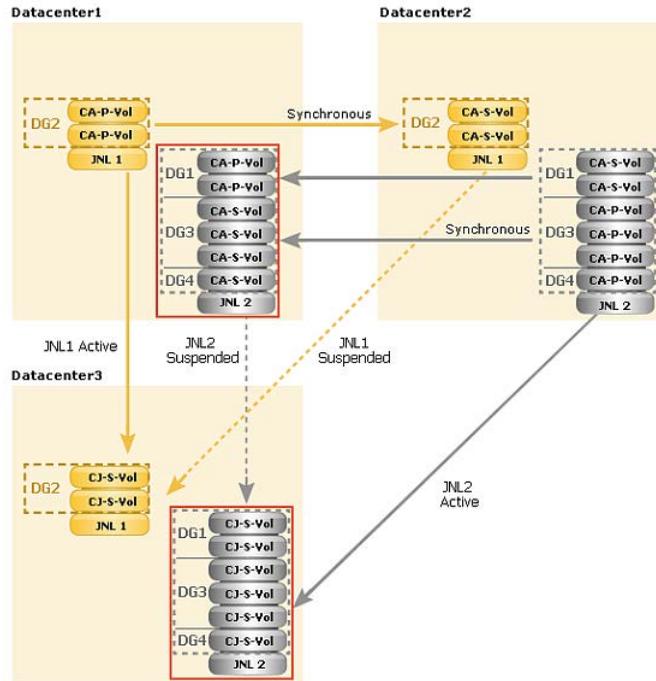
- 3 The agent runs the `horctakeover` command to switch the direction of replication.



- 4 On SYNC_TARGET, the agent runs the `paircreate` command with the `-nocopy` option for DG1.



- On PRIMARY, the agent runs the `paircreate` command with the `-nocopy` and `-nocsus` options for DG1.



This completes the migration operation without affecting any other device group. Now, DG1 replicates synchronously from Datacenter 2 to Datacenter 1, and DG1 is a part of JNL 2 for asynchronous replication.

Agent functions

Table 1-1 Agent functions

Function	Description
online	<ul style="list-style-type: none"> ■ If the value of the DefaultMode attribute of the resource is PRIMARY or SYNC_TARGET, the online function considers the local devices of the synchronous device group to identify if failover is required. If the value of the DefaultMode attribute of the resource is ASYNC_TARGET, and the value of the FailoverMode attribute is 0, the online function considers the local devices of the asynchronous device group. ■ If the state of all local devices is read-write enabled, the agent makes the devices writable by creating a lock file on the local host. ■ The agent considers P-VOL devices writable and takes no action other than going online, regardless of their status. ■ Configuration Type 1: If one or more devices are not in a writable state, the agent takes appropriate actions depending on the DefaultMode attribute value of the resource to enable read-write access to the devices. ■ Configuration Type 2: If one or more devices are not in a writable state, the agent removes the active and suspended asynchronous links from existing journal groups, performs takeover for synchronous replication link and then recreates the active and suspended links in appropriate journal groups again.
offline	<p>The agent removes the lock file that was created for the resource by the online function. The agent does not run any CCI commands because taking the resource offline is not indicative of an intention to give up the devices.</p>
monitor	<p>Verifies the existence of the lock file to determine the resource status. If the lock file exists, the agent reports the status of the resource as online. If the lock file does not exist, the agent reports the status of the resource as offline.</p> <p>For Configuration Type 1, that is, when the value of the FailoverMode attribute is zero: If the value of the DefaultMode attribute of the ONLINE resource is SYNC_TARGET, the monitor function checks the state of the asynchronous device group on PRIMARY. If the state is PSUS, the monitor function initiates resynchronization of the asynchronous link from a PRIMARY host where the VCS engine is running.</p> <p>The LinkMonitor and SplitTakeover attributes also impact the behavior of the monitor function. For more information, see the descriptions of the LinkMonitor and SplitTakeover attributes.</p>

Table 1-1 Agent functions (*continued*)

Function	Description
open	<p>Removes the lock file from the host on which this function is called. This functionality prevents potential concurrency violation if the group fails over to another node.</p> <p>Note that the agent does not remove the lock file if the agent starts after the following command:</p> <pre>hastop<-all -local> -force</pre>
clean	<p>Determines if it is safe to fault the resource if the online entry point fails or times out. The main consideration is whether a management operation was in progress when the online thread timed out and was killed. If a management operation was in progress, it could potentially leave the devices in an unusable state.</p>
info	<p>Reports the current role and status of the devices in the device group. This function can be used to verify the device state and to monitor dirty track trends.</p> <p>If the value of the DefaultMode attribute of the HTC3DC resource is PRIMARY or SYNC_TARGET, the role and status of the synchronous device group are reported.</p> <p>If the value of the DefaultMode attribute of the HTC3DC resource is ASYNC_TARGET and that of the FailoverMode attribute is zero (Configuration Type 1), the role and status of the asynchronous device group are reported.</p>

Table 1-1 Agent functions (*continued*)

Function	Description
action	<p>The agent supports the following actions using the <code>hares -action</code> command from the command line:</p> <ul style="list-style-type: none"> ■ <code>pairdisplay</code>—Displays information about all devices. Takes SYNC or ASYNC as action arguments. If the action argument is SYNC, executes the <code>pairdisplay</code> command for synchronous device group. If the action argument is ASYNC, then executes the <code>pairdisplay</code> command for the asynchronous device group. ■ <code>pairresync</code>—Takes SYNC or ASYNC as action arguments. If the action argument is SYNC, resynchronizes the S-VOL devices of the synchronous device group from the VCS command line after connectivity failures are detected and corrected. If the action argument is ASYNC, then it uses the asynchronous device group instead. ■ <code>pairresync-swaps</code>—Takes SYNC or ASYNC as action arguments. If the action argument is SYNC, promotes the S-VOLs to P-VOLs and resynchronizes the original P-VOLs of the synchronous device groups. If the action argument is ASYNC, it promotes the S-VOLs to P-VOLs and resynchronizes the original P-VOLs of the asynchronous device groups. ■ <code>localtakeover</code>—Makes the local devices write-enabled. If the action argument is SYNC, executes the <code>horctakeover</code> command for synchronous device groups. If the action argument is ASYNC, then executes the command for asynchronous device groups. If there is no action argument specified, it executes the command for sync group for SYNC_TARGET or for async group if default mode is ASYNC_TARGET. ■ <code>paircreate</code>—Creates the replication pairs. If the action argument is SYNC, it creates the synchronous replication pair for the synchronous device group. If the action argument is ASYNCPAIR, it creates asynchronous replication pair between the current site and geographically distant remote site. If the action argument is ASYNCPSUS, it creates asynchronous suspended replication pair between the current site and geographically distant remote site. ■ <code>pairsplit</code>—Executes the <code>pairsplit</code> command with the <code>-s</code> option to remove the replication link. If the action argument is SYNC, it splits the synchronous replication. If the action argument is ASYNC, it splits asynchronous replication pair.

Table 1-1 Agent functions (*continued*)

Function	Description
actionGetCurrentRPO	<p data-bbox="559 326 1209 407">Fetches the current point in time RPO. The agent performs this action function on the disaster recovery (DR) system where the ComputeDRSLA attribute is set to 1.</p> <p data-bbox="559 425 884 447">The RPO is computed in seconds.</p> <p data-bbox="559 465 1209 487">Note: The agent does not compute the RPO when the group is frozen.</p> <p data-bbox="559 505 1209 564">The agent does not store the computed RPO; make a note of the RPO for future reference.</p>

Note: The agent uses the following internal action functions to compute the RPO: StartRPOComputation, StopRPOComputation, StartWriter, ReportRPOData.

About the agent's online function

If the value of the DefaultMode attribute of the resource is PRIMARY or SYNC_TARGET, the online function considers the local devices of the synchronous device group to identify if failover is required. If the DefaultMode attribute value of the resource is ASYNC_TARGET and the FailoverMode attribute value is zero, the online function considers the local devices of the asynchronous device group.

If the state of all local devices is read-write enabled, the agent makes the devices writable by creating a lock file on the local host.

For Configuration Type 1, that is, when the value of the FailoverMode attribute of the resource is zero: If the DefaultMode attribute value of the resource is PRIMARY or SYNC_TARGET, the agent checks the device status on the VCS node with ASYNC_TARGET as default mode. If the status of the asynchronous devices is SSWS, it indicates that the failback procedure must be executed on the PRIMARY node. If the DefaultMode attribute value of the resource is SYNC_TARGET, the agent logs an error message and returns. If it is PRIMARY, then agent executes failback in the following manner:

- Removes the synchronous device group pair.
- Resynchronizes the PRIMARY devices from ASYNC_TARGET devices.
- Executes `horctakeover` on asynchronous device group to reverse the roles.
- Creates the synchronous pair, which was removed earlier.

For Configuration Type 2, that is, when the value of the FailoverMode attribute of the resource is one, the agent executes failover in the following manner:

- Runs the `pairsplit` command to remove the suspended asynchronous link between `SYNC_TARGET` and `ASYNC_TARGET`.
- Runs the `pairsplit` command to remove the active asynchronous link between `PRIMARY` and `ASYNC_TARGET`.
- Runs the `paircreate` command to create the active asynchronous link between `SYNC_TARGET` and `ASYNC_TARGET`.
- Runs the `paircreate` command to create the suspended asynchronous link between `PRIMARY` and `ASYNC_TARGET`.

In a typical failover scenario, if one or more devices are not in a writable state, the agent runs the `horctakeover` command to enable read-write access to the devices. If `horctakeover` exits with an error (exit code > 5), for example due to a timeout, then the agent flushes and freezes the group to indicate that user-intervention is required to identify the cause of the error.

When there is a failover using the synchronous device group, if the `horctakeover` command returns 225 as the exit code and the fence level is NEVER, the agent tries to issue `pairresync -swaps` for the S-VOL devices in the synchronous device group.

For S-VOL synchronous devices in any state other than SSUS, SSWS, or SMPL, depending on the type of configuration, the agent does the following:

- Configuration Type 1
The agent runs the `horctakeover` command and makes the devices writable.
- Configuration Type 2
The agent runs the `pairsplit`, `horctakeover`, and `paircreate` commands. These commands make the devices writable and rebuilds asynchronous links correctly.

For both the configuration types, the time required for failover depends on the health of the original primary and on the timeout value defined for the RAID Manager in the `horcm` configuration file for the device group.

The agent considers P-VOL synchronous devices writable and takes no action other than going online, regardless of their status.

If the S-VOL synchronous devices are in the COPY state, the agent runs the `horctakeover` command (Configuration Type 1) or the `pairsplit`, `horctakeover`, and `paircreate` commands (Configuration Type 2) after one of the following:

- The synchronization from the primary completes.
- The `OnlineTimeout` period of the function expires, in which case, the `horctakeover` command is not executed, and the resource faults. The agent validates that the value of `OnlineTimeout` for the HTC3DC type is 10 seconds

or more. If this value is set to less than 10 seconds, the agent logs an appropriate error message.

If the DefaultMode attribute value is SYNC_TARGET, and the FailoverMode attribute value is zero, after successful failover, the agent checks the status of devices belonging to asynchronous device group on PRIMARY node. If the status is PSUS, it resynchronizes the asynchronous device group by executing the `pairresync` command on PRIMARY node.

Installing and removing the agent for Hitachi/HP XP 3 Data Center Replication

This chapter includes the following topics:

- [Before you install the agent for Hitachi/HP XP 3 Data Center Replication](#)
- [Installing the agent for Hitachi/HP XP 3 Data Center Replication](#)
- [Removing the agent for Hitachi/HP XP 3 Data Center Replication](#)

Before you install the agent for Hitachi/HP XP 3 Data Center Replication

Before you install the Cluster Server agent for Hitachi/HP XP 3 Data Center Replication, ensure that you install and configure the VCS on all nodes in the cluster.

Set up replication and the required hardware infrastructure. For information about setting up Oracle RAC environment, refer to the *Storage Foundation for Oracle RAC Configuration and Upgrade Guide*.

Installing the agent for Hitachi/HP XP 3 Data Center Replication

You must install the Hitachi/HP XP 3 Data Center Replication agent on each node in the cluster. In global cluster environments, install the agent on each node in each cluster.

To install the agent in a VCS environment

- 1 Download the Agent Pack from the Symantec Operations Readiness Tools (SORT) site: <https://sort.symantec.com/agents>.
You can download the complete Agent Pack tar file or the individual agent tar file.
- 2 Uncompress the file to a temporary location, say /tmp.
- 3 If you downloaded the complete Agent Pack tar file, navigate to the directory containing the package for the platform running in your environment.

If you downloaded the individual agent tar file, navigate to the pkgs directory (for AIX, and Solaris), or the rpms directory (for Linux).

- 4 Log in as a superuser.
- 5 Install the package.

Note: On successful installation of the agent, if VCS is running, the agent types definition is automatically added to the VCS configuration.

Removing the agent for Hitachi/HP XP 3 Data Center Replication

Before you attempt to remove the agent, make sure the application service group is not online.

You must remove the HTC agent from each node in the cluster.

To remove the agent, type the following command on each node. Answer prompts accordingly:

Configuring the agent for Hitachi/HP XP 3 Data Center Replication

This chapter includes the following topics:

- [Configuration concepts for the Hitachi/HP XP 3 Data Center Replication agent](#)
- [Before you configure the agent for Hitachi/HP XP 3 Data Center Replication](#)
- [Configuring the agent for Hitachi/HP XP 3 Data Center Replication](#)

Configuration concepts for the Hitachi/HP XP 3 Data Center Replication agent

Review the resource type definition and attribute definitions for the agent.

Resource type definition for the Hitachi/HP XP 3DC agent

The resource type definition defines the agent in VCS.

```
type HTC3DC (
    static keylist RegList = { ComputeDRSLA }
    static keylist SupportedActions = { pairedisplay, pairresync,
    localtakeover, pairresync-swaps, pairsplit, paircreate,
    ReportRPOData, StartWriter, GetCurrentRPO, StartRPOComputation,
    StopRPOComputation }
    static str ArgList[] = { BaseDir, GroupName, Instance, SplitTakeover,
    LinkMonitor, AsyncGroupName, DefaultMode, FailoverMode, AsyncPairCtg,
```

```

AsyncPSusCtg, JournalPair, JournalPSus, ComputeDRSLA }
static int NumThreads = 1
static int OnlineTimeout = 600
static int MonitorInterval = 300
static int ActionTimeout = 150
str BaseDir = "/HORCM/usr/bin/"
str GroupName
int Instance
int SplitTakeover
int LinkMonitor
str AsyncGroupName
str DefaultMode = "NONE"
int FailoverMode
int AsyncPairCtg
int AsyncPSusCtg
str JournalPair{} = { Journal_Primary=0, Journal_Secondary=1 }
str JournalPSus{} = { Journal_Primary=0, Journal_Secondary=1 }
temp str TargetFrozen
temp str OldState = ""
int ComputeDRSLA
temp boolean Tagging = 0
)

```

Attribute definitions for the Hitachi/HP XP 3DC agent

Following are the attributes associated with the agent:

Table 3-1 Agent functions

Attribute	Description
BaseDir	Path to the RAID Manager Command Line interface. Type-dimension: string-scalar Default: /HORCM/usr/bin/
Instance	The Instance number of the synchronous and asynchronous device groups that the agent manages. Multiple device groups can have the same instance number. Do not define the attribute if the instance number is zero. Type-dimension: integer-scalar

Table 3-1 Agent functions (*continued*)

Attribute	Description
GroupName (Synchronous device group name)	Name of the synchronous device group that the agent manages. Type-dimension: string-scalar
FailoverMode	Defines the type of configuration. If the value of this attribute is 0, it indicates that the configuration is Configuration Type 1. If the value of this attribute is 1, it indicates that the configuration is Configuration Type 2. Type-dimension: integer-scalar
AsyncGroupName	Name of the asynchronous device group that the agent manages according to the direction of the replication. Configuration Type 1: Set this attribute on all nodes of production and DR clusters. Configuration Type 2: Set two different values for this attribute: <ul style="list-style-type: none"> ■ For all the hosts in PRIMARY: Asynchronous device group from PRIMARY to ASYNC_TARGET ■ For all the hosts in SYNC_TARGET: Asynchronous device group from SYNC_TARGET to ASYNC_TARGET Type-dimension: string-scalar
DefaultMode	This is the original mode or role of each data center. It can be either PRIMARY, SYNC_TARGET, or ASYNC_TARGET. You must set this attribute on each cluster node in production and DR clusters appropriately. Type-dimension: string-scalar Default: NONE

Table 3-1 Agent functions (*continued*)

Attribute	Description
AsyncPairCtg	<p>Applicable to Configuration Type 2; you must configure this attribute when the value of the FailoverMode attribute is 1.</p> <p>Asynchronous consistency group ID associated with the P-VOL PAIR devices that the agent manages according to the direction of the replication.</p> <p>You must specify two different values for this attribute:</p> <ul style="list-style-type: none"> ■ For all the hosts in the PRIMARY: Asynchronous consistency group ID from PRIMARY to ASYNC_TARGET (JNL 1 asynchronous active direction) ■ For all the hosts in the SYNC_TARGET: Asynchronous consistency group ID from SYNC_TARGET to ASYNC_TARGET (JNL 2 asynchronous active direction) <p>Type-dimension: integer-scalar</p>
AsyncPSusCtg	<p>Applicable to Configuration Type 2; you must configure this attribute when the value of the FailoverMode attribute is 1.</p> <p>Asynchronous consistency group ID associated with the P-VOL PSUS devices that the agent manages according to the direction of the replication.</p> <p>You must specify two different values for this attribute:</p> <ul style="list-style-type: none"> ■ For all the hosts in the PRIMARY: Asynchronous consistency group ID from PRIMARY to ASYNC_TARGET (JNL 1 asynchronous suspended direction) ■ For all the hosts in the SYNC_TARGET: Asynchronous consistency group ID from SYNC_TARGET to ASYNC_TARGET (JNL 2 asynchronous suspended direction) <p>Type-dimension: integer-scalar</p>

Table 3-1 Agent functions (*continued*)

Attribute	Description
JournalPair	<p>Applicable to Configuration Type 2; you must configure this attribute when the value of the FailoverMode attribute is 1.</p> <p>Journal ID associated with devices that the agent manages according to the direction of the replication.</p> <p>This means that will need to specify two different values identifying for each one a Journal_Primary and a Journal_Secondary:</p> <ul style="list-style-type: none"> ■ For all the hosts in the PRIMARY: Journal ID in the configuration: PRIMARY to SYNC_TARGET to ASYNC_TARGET <ul style="list-style-type: none"> ■ Journal_Primary: JNL 1 asynchronous active ■ Journal_Secondary: JNL 1 asynchronous suspended ■ For all the hosts in the SYNC_TARGET: Journal ID in the configuration: SYNC_TARGET to PRIMARY to ASYNC_TARGET <ul style="list-style-type: none"> ■ Journal_Primary: JNL 2 asynchronous active ■ Journal_Secondary: JNL 2 asynchronous suspended <p>Type-dimension: string-scalar</p>
JournalPSus	<p>Applicable to Configuration Type 2; you must configure this attribute when the value of the FailoverMode attribute is 1.</p> <p>Journal id associated to devices that the agent manages according to the direction of the replication.</p> <p>This means that will need to specify two different values identifying for each one a Journal_Primary and a Journal_Secondary.</p> <ul style="list-style-type: none"> ■ For all the hosts in the PRIMARY: Journal ID in the configuration: SYNC_TARGET to PRIMARY to ASYNC_TARGET <ul style="list-style-type: none"> ■ Journal_Primary: JNL 2 asynchronous active ■ Journal_Secondary: JNL 2 asynchronous suspended ■ For all the hosts in the SYNC_TARGET: Journal ID in the configuration: PRIMARY to SYNC_TARGET to ASYNC_TARGET <ul style="list-style-type: none"> ■ Journal_Primary: JNL 1 asynchronous active ■ Journal_Secondary: JNL 1 asynchronous suspended <p>Type-dimension: string-scalar</p>

Table 3-1 Agent functions (*continued*)

Attribute	Description
SplitTakeover	<p>A flag that determines whether the agent permits a failover to S-VOL devices if the synchronous replication link is disconnected (when S-VOL devices are in the SSUS state) or is manually suspended (when P-VOL devices are in the PSUS state).</p> <p>Type-dimension: integer-scalar</p> <p>Default: 0</p>
LinkMonitor	<p>A flag that defines whether the agent periodically attempts to resynchronize the S-VOL side if the replication link is disconnected. The agent uses the <code>pairresync</code> command to resynchronize arrays.</p> <p>The value 1 indicates that when the synchronous replication link is disconnected, the agent periodically attempts to resynchronize the S-VOL side using the <code>pairresync</code> command.</p> <p>If the value of the LinkMonitor attribute is set to 2, the agent generates SNMP traps or alerts when the status of the attached P-VOL and S-VOL changes. If the status of the configured Hitachi/HP XP 3 device changes to PSUE, PSUS, SSUS, or SSWS, the agent generates an SNMP trap indicating that the resource health has gone down. For all other types of status changes of the configured Hitachi/HP XP 3 devices, the agent generates an SNMP trap indicating that the resource health has improved. An error or information type message is logged by the agent in the VCS High Availability engine log- "The state of <i>P-VOL/S-VOL</i> devices in device group <i>device group name</i> has changed from <i>previous state</i> to <i>current state</i>."</p> <p>Setting LinkMonitor does not affect the SplitTakeover behavior. However, you can minimize the time during which the P-VOL is in the PSUE state by setting the LinkMonitor attribute.</p> <p>Type-dimension: integer-scalar</p> <p>Default: 0</p>

Table 3-1 Agent functions (*continued*)

Attribute	Description
ComputeDRSLA	<p>Used to enable or disable Recovery Point Objective (RPO) computation.</p> <p>Set this attribute on any ASYNC_TARGET node in the disaster recovery (DR) cluster. Setting this attribute to 1 starts the RPO computation process. Ensure that you reset this attribute to 0 after you use the GetCurrentRPO action function to check the RPO.</p> <p>Type-dimension: integer-scalar</p> <p>Default value : 0</p>
Tagging	An internal attribute that is used for maintaining the process of computing RPO.

About the SplitTakeover attribute for the agent

The SplitTakeover attribute determines whether the agent permits a failover to S-VOL devices if the replication link is disconnected, (that is, if P-VOL devices are in the PSUE state) or if the replication is manually suspended, if SVOL devices are in SSUS state.

SplitTakeover attribute = 0

The default value of the SplitTakeover attribute is 0.

The default value indicates that the agent does not permit a failover to S-VOL devices if the P-VOL devices are in the PSUE state, or if SVOL devices are in SSUS state. If failover occurs when the replication link is disconnected, data loss may occur because the S-VOL devices may not be in sync.

If a device group is made up of multiple devices, then, in case of a link failure, the state of each device changes on an individual basis. This change is not reflected on the device group level. Only those devices to which an application made a write after a link failure change their state to PSUE. Other devices in the same device group retain their state to PAIR.

Note: Setting LinkMonitor does not affect the SplitTakeover behavior. However you can minimize the time during which the P-VOL is in the PSUE by setting the LinkMonitor attribute.

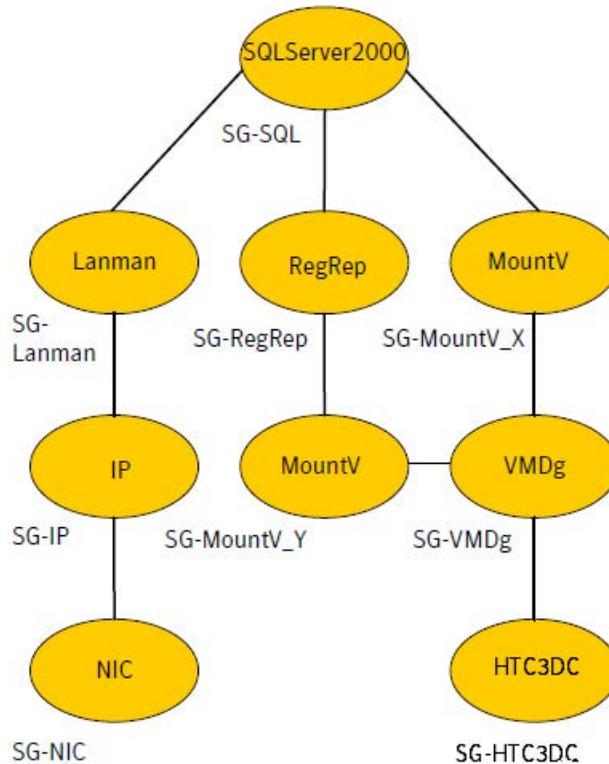
SplitTakeover attribute = 1

If the value of SplitTakeover is 1, the agent tries to make the SVOL devices writable, irrespective of the state of PVOL devices. Hence, even if there is a replication link failure, or the primary array fails, the agent attempts to failover to the S-VOL devices.

Sample configuration for the Hitachi/HP XP 3DC agent

Figure 3-1 shows a dependency graph of a VCS service group that has a resource of type HTC3DC.

Figure 3-1 VCS service group with resource type HTC3DC



You can configure a resource of type HTC3DC in the main.cf file as:

Sample configuration for Configuration Type 1

From the disaster recovery cluster:

```
HTC3DC htc3dc (
  GroupName = vg01
  Instance = 1
  AsyncGroupName = vg02
  DefaultMode = ASYNC_TARGET
)
```

From the production cluster:

```
HTC3DC HTC3DC (
  GroupName = vg01
  Instance = 1
  AsyncGroupName = vg02
  DefaultMode @PRIMARY_SYSTEM = PRIMARY
  DefaultMode @SYNC_TARGET_SYSTEM = SYNC_TARGET
)
```

Sample configuration for Configuration Type 2

```
HTC3DC htc3dc (
  GroupName = vg01
  Instance = 1
  AsyncGroupName @PRIMARY_SYS1 = dg01
  AsyncGroupName @PRIMARY_SYS2 = dg01
  AsyncGroupName @SYNC_TARGET_SYS1 = dvg01
  DefaultMode @PRIMARY_SYS1 = PRIMARY
  DefaultMode @PRIMARY_SYS2 = PRIMARY
  DefaultMode @SYNC_TARGET_SYS1 = SYNC_TARGET
  FailoverMode = 1
  AsyncPairCtg @PRIMARY_SYS1 = 6
  AsyncPairCtg @PRIMARY_SYS2 = 6
  AsyncPairCtg @SYNC_TARGET_SYS1 = 8
  AsyncPSusCtg @PRIMARY_SYS1 = 9
  AsyncPSusCtg @PRIMARY_SYS2 = 9
  AsyncPSusCtg @SYNC_TARGET_SYS1 = 7
  JournalPair @PRIMARY_SYS1 = { Journal_Primary = 0,
  Journal_Secondary = 0 }
  JournalPair @PRIMARY_SYS2 = { Journal_Primary = 0,
  Journal_Secondary = 0 }
  JournalPair @SYNC_TARGET_SYS1 = { Journal_Primary = 1,
  Journal_Secondary = 1 }
  JournalPSus @PRIMARY_SYS1 = { Journal_Primary = 1,
  Journal_Secondary = 1 }
```

```
JournalPSus @PRIMARY_SYS2 = { Journal_Primary = 1,  
Journal_Secondary = 1 }  
JournalPSus @SYNC_TARGET_SYS1 = { Journal_Primary = 0,  
Journal_Secondary = 0 }  
)
```

Before you configure the agent for Hitachi/HP XP 3 Data Center Replication

Before you configure the agent, review the following information:

- Verify that you have installed the agent on all systems in the cluster.
- Verify the hardware setup for the agent.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See “[About cluster heartbeats](#)” on page 39.

About cluster heartbeats

In a replicated data cluster, ensure robust heartbeating by using dual, dedicated networks over which the Low Latency Transport (LLT) runs. Additionally, you can configure a low-priority heartbeat across public networks.

In a global cluster, VCS sends ICMP pings over the public network between the two sites for network heartbeating. To minimize the risk of split-brain, VCS sends ICMP pings to highly available IP addresses. VCS global clusters also notify the administrators when the sites cannot communicate.

To minimize the chances of split-brain, use the steward process.

Hitachi/HP-XP arrays do not support a native heartbeating mechanism between the arrays. The arrays send a support message on detecting replication link failure. You can take appropriate action to recover from the failure and to keep the devices in a synchronized state. The Hitachi/HP XP 3DC agent supports those actions that can automate the resynchronization of devices after a replication link outage is corrected.

Configuring the agent for Hitachi/HP XP 3 Data Center Replication

You can configure clustered application in a disaster recovery environment by:

- Setting up the initial 3 DC Multi-target replication configuration between the 3 datacenters, according to the chosen configuration type.
- Synchronizing the devices
- Adding the Hitachi/HP XP 3 Data Center Replication agent to the service group

After configuration, the application service group must follow the dependency diagram.

See [Figure 3-1](#) on page 37.

Configuring the agent to compute RPO

In a global cluster environment, the agent for Hitachi/HP XP 3 Data Center Replication can compute the recovery point objective (RPO), which is a disaster recovery (DR) SLA. In a DR configuration where data is replicated asynchronously to the DR site, the DR site data is not always as current as the primary site data.

RPO is the maximum acceptable amount of data loss in case of a disaster at the primary site. The agent computes RPO in terms of time, that is, in seconds.

Before you configure the agent to compute the RPO, ensure that the following pre-requisites are met:

- The service group containing the HTC3DC resource and the VxVM disk group resource are online at the production site.
- The disk group resource is dependent on the HTC3DC resource.

To configure the agent to compute the RPO:

- 1 In the DR cluster, on any one of the nodes where devices are asynchronously replicated and where the service group is configured, run the following command to start the RPO computation:

```
hares -modify HTC3DC_resource_name ComputedDRSLA 1 -sys system_name.
```

- 2 Run the following command on the same node in the DR cluster:

```
hares -action HTC3DC_resource_name GetCurrentRPO -sys system_name
```

The action entry point displays the RPO. The agent does not store the computed RPO; make a note of the RPO for future reference.

If the RPO is not reported, it indicates that the agent needs more time to finish computing the RPO. Wait for some more time before you run the GetCurrentRPO action function again.

- 3 To stop RPO computation, run the following command:

```
hares -modify HTC3DC_resource_name ComputedDRSLA 0 -sys system_name
```

How the agent computes the replication RPO

Depending on the way 3DC is configured, the agent computes the replication RPO in the following manner for Configuration Type 1 and Configuration Type 2:

- Configuration Type 1
 - If the 3DC setup is in a multi-target mode, the agent computes the replication RPO by running the `pairdisplay` command for the ASYNC pair.
 - If the 3DC setup is in a cascaded mode, that is, the service group is online on SYNC_TARGET, then the agent computes the replication RPO by running the `pairdisplay` command with the `-FCA` option for the SYNC pair. The agent uses the `-FCA` option to get the status of the corresponding ASYNC device.
- Configuration Type 2
 - The 3DC setup is always in a multi-target mode; the agent computes the replication RPO by running the `pairdisplay` command for the ASYNC pair.

Testing VCS disaster recovery support with Hitachi/HP XP 3 Data Center Replication

This chapter includes the following topics:

- [Testing service group migration from PRIMARY to SYNC_TARGET](#)
- [Failure scenarios at PRIMARY](#)
- [Production cluster failure](#)
- [Performing failback from ASYNC_TARGET](#)

Testing service group migration from PRIMARY to SYNC_TARGET

After you configure the VCS agent for Hitachi/HP XP 3 Data Center Replication, verify that the application service group that contains the HTC3DC resource can migrate to hosts at SYNC_TARGET.

To test the service group migration in a VCS extended cluster or an RDC

- 1 Fail over the application service group that contains the HTC3DC resource from the PRIMARY to the SYNC_TARGET.

Run the following command to switch the application service group from the PRIMARY node to the SYNC_TARGET node in the Replicated Data Cluster (RDC) or the VCS extended cluster.

```
# hagr -switch app_group -sys sync_target_system
```

VCS brings the application service group online on the SYNC_TARGET node.

- 2 Verify that the devices that are managed by the HTC3DC resource at the SYNC_TARGET node are write-enabled.
The HTC3DC resource comes online.
- 3 Verify that the state of the synchronous device group is PAIR.
- 4 For Configuration Type 1: Verify that asynchronous replication link which was suspended due to a role change in the synchronous device group is resumed.
For Configuration Type 2: Verify that the asynchronous device group managed by the HTC3DC resource on SYNC_TARGET is now added to the active journal 2 between SYNC_TARGET and ASYNC_TARGET.
- 5 For Configuration Type 1: Verify that the original Multi-target configuration has now changed to Cascaded configuration, where the direction of replication is from SYNC_TARGET to PRIMARY to ASYNC_TARGET.
For Configuration Type 2: Verify that the asynchronous device group managed by the HTC3DC resource on PRIMARY is now added to the suspended journal 2.

Failure scenarios at PRIMARY

[Table 4-1](#) lists the failure scenarios at the PRIMARY site and describes the behavior of VCS and the agent in response to the failure. This table also describes the steps you can take to verify disaster recovery.

Table 4-1 Failure scenarios at the PRIMARY site with VCS agent for Hitachi/HP XP 3DC

Failure scenario	Description and VCS response
Application failure at PRIMARY	<p>Configuration Type 1</p> <p>An application failure at the PRIMARY causes the application service group to fault. The application service group fails over to the SYNC_TARGET node in the VCS extended cluster (or RDC).</p> <p>VCS action:</p> <p>In the RDC environment, VCS does the following based on the AutoFailOver attribute for the faulted service group:</p> <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at SYNC_TARGET. ■ 2—You must bring the service group online at SYNC_TARGET. <p>Agent response:</p> <ul style="list-style-type: none"> ■ The agent write-enables the devices at the SYNC_TARGET site. ■ The agent does the following: <ul style="list-style-type: none"> ■ Does a role reversal S-VOL to P-VOL of each device managed by HTC3DC resource in the synchronous device group. ■ Starts replication from P-VOL devices on the SYNC_TARGET to the S-VOL devices at the PRIMARY. ■ Resynchronises the asynchronous device group from PRIMARY. <p>Verification steps:</p> <ul style="list-style-type: none"> ■ Verify that the devices at SYNC_TARGET became write-enabled and the state is PAIR. ■ Verify that the state of the synchronous device group is PAIR. ■ Verify that the asynchronous replication link which was suspended due to role change in synchronous device group is resumed. ■ Verify that the original Multi-target configuration is now changed to Cascaded configuration where the replication direction is from SYNC_TARGET to PRIMARY to ASYNC_TARGET.

Table 4-1 Failure scenarios at the PRIMARY site with VCS agent for Hitachi/HP XP 3DC (*continued*)

Failure scenario	Description and VCS response
	<p>Configuration Type 2</p> <p>An application failure at the PRIMARY causes the application service group to fault. The application service group fails over to the SYNC_TARGET node in the VCS extended cluster (or RDC).</p> <p>VCS action</p> <p>In the RDC environment, VCS does the following based on the AutoFailOver attribute for the faulted service group:</p> <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at SYNC_TARGET. ■ 2—You must bring the service group online at SYNC_TARGET. <p>Agent response</p> <p>The agent does the following:</p> <ul style="list-style-type: none"> ■ Splits the active asynchronous link between PRIMARY and ASYNC_TARGET of the asynchronous group managed by the HTC3DC resource. ■ Splits the suspended asynchronous link between SYNC_TARGET and ASYNC_TARGET of the asynchronous group managed by the HTC3DC resource. ■ Does a role reversal S-VOL to P-VOL of each device managed by the HTC3DC resource in the synchronous device group. ■ Creates an asynchronous pair between SYNC_TARGET and ASYNC_TARGET using active journal 2. ■ Creates an asynchronous suspended pair between PRIMARY and ASYNC_TARGET using suspended journal 2.

Table 4-1 Failure scenarios at the PRIMARY site with VCS agent for Hitachi/HP XP 3DC (continued)

Failure scenario	Description and VCS response
<p>Storage failure at PRIMARY</p> <p>Note: This failure scenario is only applicable to Configuration Type 1.</p>	<p>Array failure at the PRIMARY node causes the service group to fault. The service group fails over to the SYNC_TARGET node in the VCS extended cluster (or RDC).</p> <p>VCS action:</p> <p>In the RDC environment, VCS does the following based on the AutoFailOver attribute for the faulted service group:</p> <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at SYNC_TARGET. ■ 2—You must bring the service group online at SYNC_TARGET. <p>Agent response:</p> <p>The agent does the following based on the SplitTakeover attribute of the HTC3DC resource:</p> <ul style="list-style-type: none"> ■ 1—The agent issues the <code>horctakeover</code> command to make the HTC3DC devices at SYNC_TARGET write-enabled. The S-VOL devices go into the SSWS state. ■ 0—The agent faults the HTC3DC resource. <p>Verification steps:</p> <ul style="list-style-type: none"> ■ If the value of the SplitTakeOver attribute is 1, verify that the devices at SYNC_TARGET are now write-enabled and that the status of the synchronous device group is SSWS. ■ The storage at PRIMARY is not available and hence, the mode of configuration is neither 3DC Multi-target nor 3DC Cascaded.

Table 4-1 Failure scenarios at the PRIMARY site with VCS agent for Hitachi/HP XP 3DC (*continued*)

Failure scenario	Description and VCS response
<p>Host failure at PRIMARY</p> <p>Note: This failure scenario is only applicable to Configuration Type 1.</p>	<p>All hosts at the PRIMARY site fail.</p> <p>Note: This failure scenario is not applicable to Configuration Type 2.</p> <p>Host failure at the PRIMARY causes the service group to fault. The service group fails over to the SYNC_TARGET node in the VCS extended cluster (or RDC).</p> <p>VCS action:</p> <p>In the RDC environment, VCS does the following based on the AutoFailOver attribute for the faulted service group:</p> <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at SYNC_TARGET. ■ 2—You must bring the service group online at the SYNC_TARGET. <p>Agent response:</p> <ul style="list-style-type: none"> ■ The agent write-enables the devices at the SYNC_TARGET site. ■ The agent does the following: <ul style="list-style-type: none"> ■ Does a role reversal S-VOL to P-VOL of each device managed by HTC3DC resource in the synchronous device group. ■ Starts replication from P-VOL devices on the SYNC_TARGET to the S-VOL devices at the PRIMARY. ■ Resynchronises the asynchronous device group from PRIMARY. <p>Verification steps</p> <ul style="list-style-type: none"> ■ Verify that the devices at SYNC_TARGET became write-enabled and the state is PAIR. ■ Verify the status of the asynchronous replication link once the PRIMARY is up and running. ■ Verify that the asynchronous replication link which was suspended due to role change in synchronous device group is resumed. ■ Verify the original Multi-target configuration is now changed to Cascaded configuration once the PRIMARY node is up and running. In the Cascaded configuration, the replication direction is from SYNC_TARGET to PRIMARY to ASYNC_TARGET.

Table 4-1 Failure scenarios at the PRIMARY site with VCS agent for Hitachi/HP XP 3DC (continued)

Failure scenario	Description and VCS response
Replication link failure at PRIMARY	<p>The replication link between the PRIMARY and SYNC_TARGET arrays fails.</p> <p>For both the configuration types, when there is a replication link failure at the PRIMARY, the following links are monitored:</p> <ul style="list-style-type: none"> ■ If the resource is online on PRIMARY or SYNC_TARGET, then the synchronous link is monitored. ■ If the resource is online at ASYNC_TARGET, then the asynchronous link is monitored. <p>VCS action: No action.</p> <p>Agent response: The Monitor function detects the replication link failure and takes the following action depending on the value of the LinkMonitor attribute:</p> <ul style="list-style-type: none"> ■ If LinkMonitor is set to 0, the agent takes no action. ■ If LinkMonitor is set to 1, when the replication link is disconnected, the agent periodically attempts to resynchronize synchronous device group S-VOL devices using the <code>pairresync</code> command. <p>In case of a replication link failure, failover to SYNC_TARGET depends on the value of the SplitTakeover attribute at SYNC_TARGET:</p> <ul style="list-style-type: none"> ■ 1—The agent issues the <code>horctakeover</code> command to make the HTC3DC devices at SYNC_TARGET write-enabled. The S-VOL devices go into the SSWS state. ■ 0—The agent faults the HTC3DC resource. <p>Verification steps: If the value of the LinkMonitor attribute is not set to 1, you must manually resynchronize the HTC3DC devices after the link is restored.</p> <p>To manually resynchronize the HTC3DC devices after the link is restored:</p> <ul style="list-style-type: none"> ■ You must initiate resync of S-VOL device using the agent's <code>pairresync</code> action where the HTC3DC resource is online. ■ In case of failover to SYNC_TARGET, verify that the devices are write-enabled and the synchronous device group state is SSWS.

Production cluster failure

Production cluster failure indicates that all hosts and the storage at PRIMARY and SYNC_TARGET are down.

This failure scenario is only applicable to Configuration Type 1.

VCS action: VCS does the following based on the ClusterFailOverPolicy global service group attribute:

- Auto—VCS automatically brings the faulted application service group online at ASYNC_TARGET.
- Manual or Connected—No action. You must bring the application service group online at ASYNC_TARGET.

Agent response:

- The agent uses the asynchronous device group for failover.
- The agent does the following at ASYNC_TARGET, based on the value of the SplitTakeOver attribute of the HTC3DC resource:
 - 1—The agent issues the `horctakeover` command to make the HTC3DC devices write-enabled. The HTC3DC devices go into the SSWS (Suspend for Swapping with S-VOL side only) state.
 - 0—The agent takes no action and the resource faults.

Verification steps

- Verify that the devices are write enabled at the ASYNC target node and their status is SSWS.
The configuration is neither in 3DC Multi-target mode nor in 3DC Cascaded mode since the storage at PRIMARY is not available.

Performing failback from ASYNC_TARGET

All the hosts and storage in the production cluster recover from failure and then initiate a failback by switching the application service group to the PRIMARY node.

Failback from ASYNC_TARGET is only applicable to Configuration Type 1.

The agent takes the following action:

- splits the SYNC replication link at SYNC_TARGET.
- resynchronizes all the data that was written at ASYNC_TARGET when the production cluster was down.
- makes all devices write-enabled at the PRIMARY.
- starts the replication to SYNC_TARGET.

The configuration is brought back to 3DC Multi-target configuration.

Index

A

- Agent functions
 - action 22
 - clean 22
 - GetCurrentRPO 22
 - info 22
 - monitor 22
 - offline 22
 - online 22
- Attributes 31
 - about SplitTakeover 36

C

- cluster
 - heartbeats 39
- Configuration Type 1
 - about 9
 - Cascaded 14
 - failover 14
 - typical setup 14
- Configuration Type 2
 - about 10
 - device group migration 16
 - typical setup 15
- configuring
 - before 30
 - samples 37

F

- Failover
 - in Configuration Type 1 14
 - in Configuration Type 2 16

O

- online function
 - about 25

R

- Recovery Point Objective (RPO)
 - Configuring RPO computation support 40

S

- sample configuration 37

T

- typical setup 13