

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

Solaris

7.0

VERITASTM

Veritas InfoScale™ Availability Agents

Last updated: 2019-01-10

Legal Notice

Copyright © 2019 Veritas Technologies LLC. All rights reserved.

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

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

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

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

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

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

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

<http://www.veritas.com>

Technical Support

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

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

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

<https://my.veritas.com>

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

Worldwide (except Japan)

CustomerCare@veritas.com

Japan

CustomerCare_Japan@veritas.com

Documentation

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

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

Documentation feedback

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

xyz@veritas.com

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

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

Veritas Services and Operations Readiness Tools (SORT)

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

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

Contents

Chapter 1	Introducing the agent for Hitachi/HP XP 3 Data Center Asynchronous Replication	6
	About the agent for Hitachi/HP XP 3 Data Center Asynchronous Replication	6
	Supported software	8
	Supported hardware for Hitachi/HP XP 3 Data Center Asynchronous Replication	8
	Typical Hitachi/HP XP 3 Data Center Asynchronous Replication setup in a VCS cluster	8
	Supported configuration	9
	Failover of a service in supported configuration	10
	Agent functions	11
	About the agent's online function	13
Chapter 2	Installing and removing the agent for Hitachi/HP XP 3 Data Center Asynchronous Replication	15
	Before you install the agent for Hitachi/HP XP 3 Data Center Asynchronous Replication	15
	Installing the agent for Hitachi/HP XP 3 Data Center Asynchronous Replication	16
	Removing the agent for Hitachi/HP XP 3 Data Center Asynchronous Replication	16
Chapter 3	Configuring the agent for Hitachi/HP XP 3 Data Center Asynchronous Replication	18
	Configuration concepts for the Hitachi/HP XP 3 Data Center Asynchronous Replication agent	18
	Resource type definition for the Hitachi/HP XP 3DC Aysnc agent	18
	Attribute definitions for the Hitachi/HP XP 3 Data Center Asynchronous Replication agent	19
	About the SplitTakeover attribute of the agent	22

Sample configuration for the Hitachi/HP XP 3DC Aysnc agent	22
Sample configurations for the Hitachi/HP XP 3DC Aysnc agent in a GCO environment	24
Before you configure the agent for Hitachi/HP XP 3 Data Center	
Asynchronous Replication	33
About cluster heartbeats	33
Configuring the agent for Hitachi/HP XP 3 Data Center Asynchronous	
Replication	33
Configuring the agent to compute RPO	34
How the agent computes the replication RPO	35

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

This chapter includes the following topics:

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

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

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

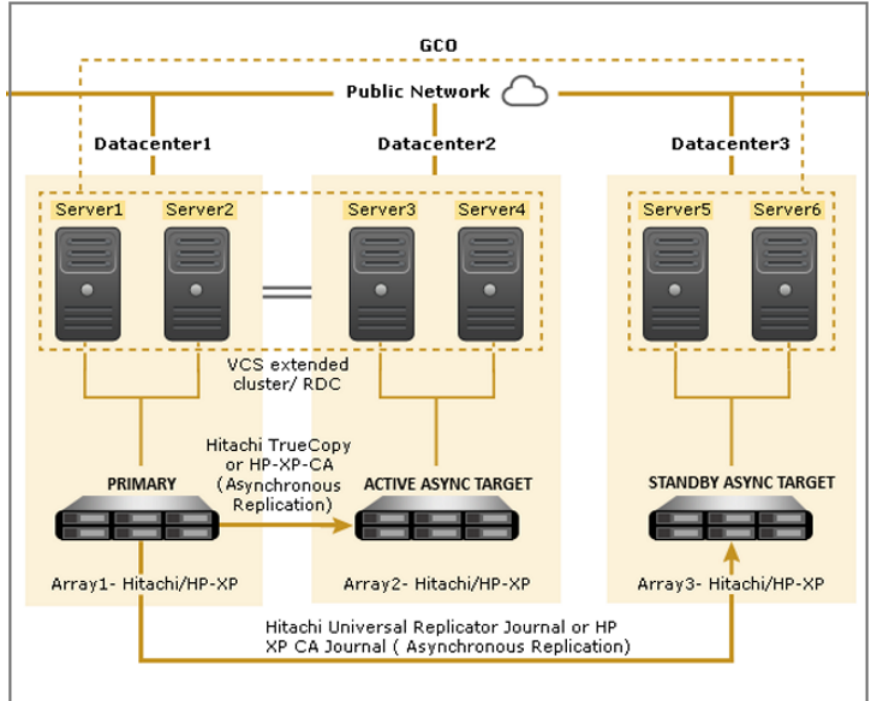
The agent monitors and manages the state of replicated Hitachi/HP XP 3 Data Center Asynchronous Replication devices that are attached to VCS nodes. The agent ensures that the system that has the Hitachi/HP XP 3 Data Center Asynchronous Replication resource online also has safe and exclusive access to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

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).
- **ACTIVE_ASYNC_TARGET:** The local datacenter to which the production data is replicated asynchronously. This is typically a part of a VCS extended cluster or a VCS RDC.
- **STANDBY_ASYNC_TARGET:** The datacenter at the DR site to which the production data is replicated asynchronously.

The Cluster Server agent for Hitachi/HP XP 3 Data Center Asynchronous Replication supports the 3 Data Center (3DC) Multi-target implementation in the following configuration:

Figure 1-1 Supported multi-target mode of operation



In this 3DC multi-target implementation, data from source is replicated asynchronously to a nearby secondary site. 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.

Supported software

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

Supported hardware for Hitachi/HP XP 3 Data Center Asynchronous Replication

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

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

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

Clustering in a Hitachi/HP XP 3DC Aysnc 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 HUR Journal P-VOL devices.
- Datacenter 2 (ACTIVE_ASYNC_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 S-VOL devices. The S-VOL

devices are paired with the P-VOL devices in the asynchronous device group, in Array1.

- Datacenter 3 (STANDBY_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.
See [“About cluster heartbeats”](#) on page 33.
- In a replicated data cluster environment, comprising Datacenter 1 (PRIMARY) and Datacenter 2 (ACTIVE_ASYNC_TARGET), all hosts are part of the same cluster. You must connect them with the dual and dedicated networks that support LLT.

Supported configuration

This graphic in a previous section displays a typical cluster setup in a Hitachi/HP XP 3DC Aysnc environment:

See [Figure 1-1](#) on page 7.

In such a typical cluster setup:

- PRIMARY and ACTIVE_ASYNC_TARGET form a VCS extended cluster or VCS RDC.
- VCS Global Cluster Option (GCO) is configured between the production cluster (includes PRIMARY and ACTIVE_ASYNC_TARGET) and the DR cluster (includes STANDBY_ASYNC_TARGET).
- In a global cluster environment at the Disaster-Recovery (DR) site, that is, Datacenter 3 (STANDBY_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 ACTIVE_ASYNC_TARGET and STANDBY_ASYNC_TARGET, for example, a link to leverage the 3DC Delta Resync feature.

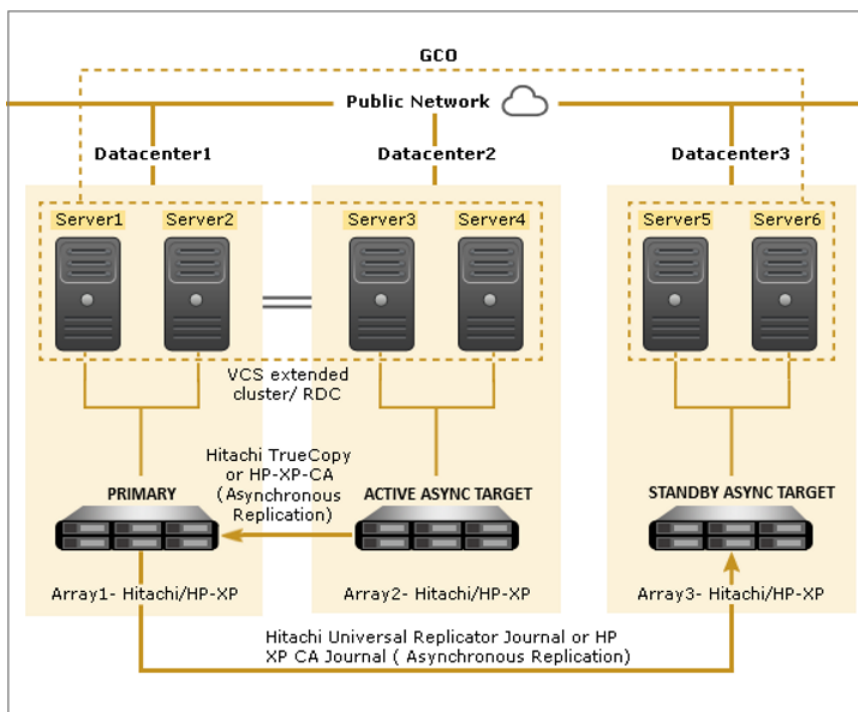
Failover of a service in supported configuration

In a 3DC Multi-target mode of operation in Configuration Type 1, production data from the PRIMARY is replicated asynchronously to ACTIVE_ASYNC_TARGET, which is part of the replicated data cluster. Simultaneously, production data from the PRIMARY is replicated asynchronously to the remote STANDBY_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 ACTIVE_ASYNC_TARGET. If there is a disaster at the production site, the application is failed over to the DR site, that is, STANDBY_ASYNC_TARGET. In this case, the agent initiates a takeover procedure on STANDBY_ASYNC_TARGET. When the production cluster is up again, the replication status needs to be brought into PAIR mode before proceeding to fail back to the original PRIMARY site.

The following graphic shows how the configuration looks when failover is successfully accomplished in case of a planned migration from the production cluster.

Figure 1-2 3DC Cascaded mode of operation



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

Agent functions

The agent performs the following functions when managing the asynchronous replication with Hitachi/HP-XP:

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 ACTIVE_ASYNC_TARGET, the online function considers the local devices of the local asynchronous device group to identify if failover is required. If the value of the DefaultMode attribute of the resource is STANDBY_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.■ 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. <p>See “About the agent's online function” on page 13.</p>
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>

Table 1-1 Agent functions (*continued*)

Function	Description
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>When the value of the FailoverMode attribute is zero: If the value of the DefaultMode attribute of the ONLINE resource is ACTIVE_ASYNC_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. See “Attribute definitions for the Hitachi/HP XP 3 Data Center Asynchronous Replication agent” on page 19.</p>
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 HTC3DCASync resource is PRIMARY or ACTIVE_ASYNC_TARGET, the role and status of the local asynchronous device group are reported.</p> <p>If the value of the DefaultMode attribute of the HTC3DCASync resource is STANDBY_ASYNC_TARGET and that of the FailoverMode attribute is zero, the role and status of the remote asynchronous device group are reported.</p>

The agent supports the following actions using the `hares -action` command:

Table 1-2 Action entry points

Function	Description
<code>pairedisplay</code>	Displays information about all devices. Takes SYNC or ASYNC as action arguments. If the action argument is SYNC, executes the <code>pairedisplay</code> command for local asynchronous device group. If the action argument is ASYNC, then executes the <code>pairedisplay</code> command for the DR 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 local asynchronous device group from the VCS command line after connectivity failures are detected and corrected. If the action argument is ASYNC, then it uses the DR asynchronous device group instead.
<code>GetCurrentRPO</code>	<p>Fetches the current point in time RPO. The agent performs this action function on the disaster recovery (DR) system where the <code>ComputedDRSLA</code> attribute is set to 1.</p> <p>The RPO is computed in seconds.</p> <p>Note: The agent does not compute the RPO when the group is frozen.</p> <p>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 `ACTIVE_ASYNC_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 `STANDBY_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.

When the value of the `FailoverMode` attribute of the resource is zero:

- If the `DefaultMode` attribute value of the resource is `PRIMARY` or `ACTIVE_ASYNC_TARGET`, the agent checks the device status on the VCS node with `STANDBY_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 `ACTIVE_ASYNC_TARGET`, the agent logs an error message and returns.
- If it is `PRIMARY`, then you need to bring the devices back into the correct states to perform failback.

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.

For S-VOL synchronous devices in any state other than `SSUS`, `SSWS`, or `SMPL`, the agent runs the `horctakeover` command and makes the devices writable.

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 asynchronous 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 after one of the following events:

- 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 `HTC3DCASync` 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 `ACTIVE_ASYNC_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 remote asynchronous device group by executing the `pairresync` command on `PRIMARY` node.

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

This chapter includes the following topics:

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

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

Before you install the VCS agent for Hitachi/HP XP 3 Data Center Asynchronous Replication, ensure that you install and configure 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 Asynchronous Replication

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

These instructions assume that you have already installed VCS.

To install the agent IPS package on an Oracle Solaris 11 system

- 1 Copy the `VRTSvcstc3dc.p5p` package from the `pkgs` directory to the system in the `/tmp/install` directory.

- 2 Disable the publishers that are not reachable as package install may fail, if any of the already added repositories are unreachable.

```
# pkg set-publisher --disable <publisher name>
```

where the publisher name is obtained using the `pkg publisher` command.

- 3 Add a file-based repository in the system.

```
# pkg set-publisher -g /tmp/install/VRTSvcstc3dc.p5p Veritas
```

- 4 Install the package.

```
# pkg install --accept VRTSvcstc3dc
```

- 5 Remove the publisher from the system.

```
# pkg unset-publisher Veritas
```

- 6 Enable the publishers that were disabled earlier.

```
# pkg set-publisher --enable publisherName
```

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 Asynchronous Replication

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

You must remove the Hitachi/HP XP 3DC Aysnc agent from each node in the cluster.

To remove the agent

- ◆ Type the following command on each node, and answer the prompts as appropriate:

```
# pkg uninstall VRTSvcstc3dc
```

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

This chapter includes the following topics:

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

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

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

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

The resource type definition defines the agent in VCS.

```
type HTC3DCASYNC (
    static keylist RegList = { ComputedRSLSA }
    static keylist SupportedActions = { pairedisplay, pairresync }
    static str ArgList[] = { BaseDir, ActiveAsyncGroupName,
```

```

        Instance, SplitTakeover, LinkMonitor, StandbyAsyncGroupName,
        DefaultMode, ComputeDRSLA }
static int NumThreads = 1
static int OnlineTimeout = 600
static int MonitorInterval = 300
static int ActionTimeout = 150
str BaseDir = "/HORCM/usr/bin/"
str ActiveAsyncGroupName
int Instance
int SplitTakeover
int LinkMonitor
str StandbyAsyncGroupName
str DefaultMode = "NONE"
temp str TargetFrozen
temp str OldState = ""
int ComputeDRSLA
temp boolean Tagging = 0
    )

```

Attribute definitions for the Hitachi/HP XP 3 Data Center Asynchronous Replication agent

Review the agent attributes before you configure a service group for HTC3DCASync.

Table 3-1 Attribute definitions

Attribute	Description
Name: BaseDir Type: String Dimension: Scalar	Path to the RAID Manager Command Line interface. Default: /HORCM/usr/bin/
Name: Instance Type: Integer Dimension: Scalar	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.
Name: ActiveAsyncGroupName Type: String Dimension: Scalar	Name of the local asynchronous device group that the agent manages.

Table 3-1 Attribute definitions (*continued*)

Attribute	Description
Name: StandbyAsyncGroupName Type: String Dimension: Scalar	Name of the DR asynchronous device group that the agent manages according to the direction of the replication. Set this attribute on all nodes of production and DR clusters.
Name: DefaultMode Type: String Dimension: Scalar	This is the original mode or role of each data center. It can be either PRIMARY, ACTIVE_ASYNC_TARGET, or STANDBY_ASYNC_TARGET. You must set this attribute on each cluster node in production and DR clusters appropriately. Default: (none)
Name: SplitTakeover Type: Integer Dimension: Scalar	A flag that determines whether the agent permits a failover to S-VOL devices if the local asynchronous 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). Default: 0

Table 3-1 Attribute definitions (*continued*)

Attribute	Description
Name: LinkMonitor Type: Integer Dimension: Scalar	<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 active asynchronous 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>Default: 0</p>
Name: ComputeDRSLA Type: Integer Dimension: Scalar	<p>Used to enable or disable Recovery Point Objective (RPO) computation.</p> <p>Set this attribute on any STANDBY_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>Default: 0</p>
Name: Tagging	<p>An internal attribute that is used for maintaining the process of computing RPO.</p>

About the SplitTakeover attribute of 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.

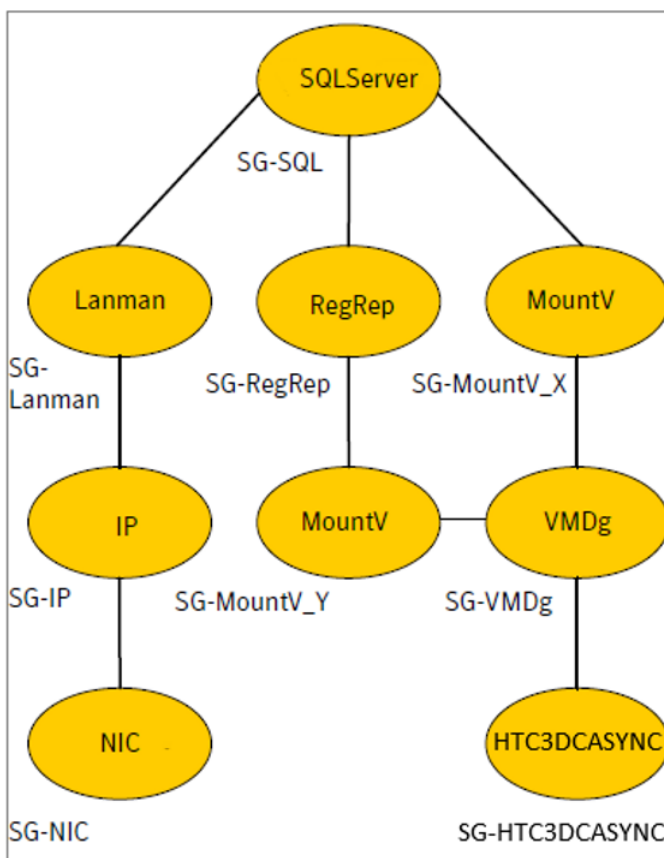
SplitTakeover attribute = 1

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 Aysnc agent

The following graphic shows a dependency graph of a VCS service group that has a resource of type HTC3DCASync.

Figure 3-1 Sample dependency graph



You can configure a resource of type HTC3DCASync in the `main.cf` file as follows:

From the disaster recovery cluster:

```

HTC3DCASync htc3dcasync (
    ActiveAsyncGroupName = vg01
    Instance = 1
    StandbyAsyncGroupName = vg02
    DefaultMode = STANDBY_ASYNC_TARGET
)
  
```

From the production cluster:

```

HTC3DCASync HTC3DCASync (
    ActiveAsyncGroupName = vg01
  )
  
```

```

Instance = 1
StandbyAsyncGroupName = vg02
DefaultMode @PRIMARY_SYSTEM = PRIMARY
DefaultMode @ACTIVE_ASYNC_TARGET_SYSTEM = ACTIVE_ASYNC_TARGET
)

```

Sample configurations for the Hitachi/HP XP 3DC Aysnc agent in a GCO environment

A sample `main.cf` file that includes a resource of the HTC3DCASYNC type at PRIMARY is as follows:

```

include "OracleASMTypes.cf"
include "types.cf"
include "CRSResource.cf"
include "CSSD.cf"
include "Db2udbTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "SybaseTypes.cf"

cluster htc3dc1 (
    UserNames = { admin = aHiaHChEIdIigQIcHF }
    ClusterAddress = "10.209.112.7"
    Administrators = { admin }
)

remotecluster htc3dc2 (
    ClusterAddress = "10.209.112.8"
)

remotecluster htc3dc3 (
    ClusterAddress = "10.209.112.9"
)

heartbeat Icmp (
    ClusterList = { htc3dc2, htc3dc3 }
    Arguments @htc3dc2 = { "10.209.112.8" }
    Arguments @htc3dc3 = { "10.209.112.9" }
)

system Host1 (

```



```

    )

system Host2 (
)

group ClusterService (
    SystemList = { Host1 = 0, Host2 = 1 }
    AutoStartList = { Host1, Host2 }
    OnlineRetryLimit = 3
    OnlineRetryInterval = 120
)

Application wac (
    StartProgram = "/opt/VRTSvcs/bin/wacstart"
    StopProgram = "/opt/VRTSvcs/bin/wacstop"
    MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
    RestartLimit = 3
)

IP gcoip (
    Device = net0
    Address = "10.209.112.7"
    NetMask = "255.255.252.0"
)

NIC gconic (
    Device = net0
)

gcoip requires gconic
wac requires gcoip

// resource dependency tree
//
// group ClusterService
// {
//   Application wac
//   {
//     IP gcoip
//     {
//       NIC gconic
//     }
//   }
// }

```

```
//      }
//  }

group htc3dcsg (
    SystemList = { Host1 = 0, Host2 = 1 }
    ClusterList = { htc3dc3 = 2, htc3dc2 = 1, htc3dc1 = 0 }
    AutoStartList = { Host1, Host2 }
)

DiskGroup dg_res (
    DiskGroup = 3dctestdg
)

FileOnOff app_res (
    PathName = "/3dcmount/a.txt"
)

HTC3DCASync htc3dc_res (
    ActiveAsyncGroupName = rep_aa
    Instance = 2
    StandbyAsyncGroupName = rep_sa
    DefaultMode @Host1 = PRIMARY
    DefaultMode @Host2 = PRIMARY
)

Mount mount_res (
    MountPoint = "/3dcmount"
    BlockDevice = "/dev/vx/dsk/3dctestdg/3dctestvol"
    FSType = vxfs
    FsckOpt = "-y"
)

app_res requires mount_res
dg_res requires htc3dc_res
mount_res requires dg_res

// resource dependency tree
//
// group htc3dcsg
// {
// FileOnOff app_res
```

```
//      {
//      Mount mount_res
//      {
//      DiskGroup dg_res
//      {
//      HTC3DCASync htc3dc_res
//      }
//      }
//      }
//      }
```

A sample `main.cf` file that includes a resource of the HTC3DCASync type at ACTIVE_ASYNC_TARGET site is as follows:

```
include "OracleASMTypes.cf"
include "types.cf"
include "CRSResource.cf"
include "CSSD.cf"
include "Db2udbTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "SybaseTypes.cf"

cluster htc3dc2 (
  UserNames = { admin = GLMeLG1IMhMMkUMg1J }
  ClusterAddress = "10.209.112.8"
  Administrators = { admin }
)

remotecluster htc3dc1 (
  ClusterAddress = "10.209.112.7"
)

remotecluster htc3dc3 (
  ClusterAddress = "10.209.112.9"
)

heartbeat Icmp (
  ClusterList = { htc3dc1, htc3dc3 }
  Arguments @htc3dc1 = { "10.209.112.7" }
  Arguments @htc3dc3 = { "10.209.112.9" }
)
```

```

system Host3 (
)

group ClusterService (
  SystemList = { Host3 = 0 }
  AutoStartList = { Host3 }
  OnlineRetryLimit = 3
  OnlineRetryInterval = 120
)

Application wac (
  StartProgram = "/opt/VRTSvcs/bin/wacstart"
  StopProgram = "/opt/VRTSvcs/bin/wacstop"
  MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
  RestartLimit = 3
)

IP gcoip (
  Device = net0
  Address = "10.209.112.8"
  NetMask = "255.255.252.0"
)

NIC gconic (
  Device = net0
)

gcoip requires gconic
wac requires gcoip

// resource dependency tree
//
// group ClusterService
// {
//   Application wac
//     {
//       IP gcoip
//         {
//           NIC gconic
//         }
//     }
// }
// }
```

```

group htc3dcsg (
    SystemList = { Host3 = 0 }
    ClusterList = { htc3dc2 = 1, htc3dc1 = 0, htc3dc3 = 2 }
    AutoStartList = { Host3 }
)

DiskGroup dg_res (
    DiskGroup = 3dctestdg
)

FileOnOff app_res (
    PathName = "/3dcmount/a.txt"
)

HTC3DCASync htc3dc_res (
    ActiveAsyncGroupName = rep_aa
    Instance = 2
    StandbyAsyncGroupName = rep_sa
    DefaultMode = ACTIVE_ASYNC_TARGET
)

Mount mount_res (
    MountPoint = "/3dcmount"
    BlockDevice = "/dev/vx/dsk/3dctestdg/3dctestvol"
    FSType = vxfs
    FsckOpt = "-y"
)

app_res requires mount_res
dg_res requires htc3dc_res
mount_res requires dg_res

// resource dependency tree
//
// group htc3dcsg
// {
//   FileOnOff app_res
//   {
//     Mount mount_res
//     {

```

```
//      DiskGroup dg_res
//      {
//          HTC3DCASync htc3dc_res
//      }
//      }
//  }
// }
```

A sample `main.cf` file that includes a resource of the HTC3DCASync type at STANDBY_ASYNC_TARGET is as follows:

```
include "OracleASMTypes.cf"
include "types.cf"
include "CRSResource.cf"
include "CSSD.cf"
include "Db2udbTypes.cf"
include "MultiPrivNIC.cf"
include "OracleTypes.cf"
include "PrivNIC.cf"
include "SybaseTypes.cf"

cluster htc3dc3 (
    UserNames = { admin = cmnFmhMjnInnLvnHmk }
    ClusterAddress = "10.209.112.9"
    Administrators = { admin }
)

remotecluster htc3dc1 (
    ClusterAddress = "10.209.112.7"
)

remotecluster htc3dc2 (
    ClusterAddress = "10.209.112.8"
)

heartbeat Icmp (
    ClusterList = { htc3dc1, htc3dc2 }
    Arguments @htc3dc1 = { "10.209.112.7" }
    Arguments @htc3dc2 = { "10.209.112.8" }
)

system Host4 (
)
```

```

group ClusterService (
  SystemList = { Host4 = 0 }
  AutoStartList = { Host4 }
  OnlineRetryLimit = 3
  OnlineRetryInterval = 120
)

Application wac (
  StartProgram = "/opt/VRTSvcs/bin/wacstart"
  StopProgram = "/opt/VRTSvcs/bin/wacstop"
  MonitorProcesses = { "/opt/VRTSvcs/bin/wac" }
  RestartLimit = 3
)

IP gcoip (
  Device = net0
  Address = "10.209.112.9"
  NetMask = "255.255.252.0"
)

NIC gconic (
  Device = net0
)

gcoip requires gconic
wac requires gcoip

// resource dependency tree
//
// group ClusterService
// {
//   Application wac
//     {
//       IP gcoip
//         {
//           NIC gconic
//         }
//     }
// }

group htc3dcs (

```

```

SystemList = { Host4 = 0 }
ClusterList = { htc3dc2 = 1, htc3dc1 = 0, htc3dc3 = 2 }
Authority = 1
AutoStartList = { Host4 }
)

DiskGroup dg_res (
    DiskGroup = 3dctestdg
)

FileOnOff app_res (
    PathName = "/3dcmount/a.txt"
)

HTC3DCASync htc3dc_res (
    ActiveAsyncGroupName = rep_aa
    Instance = 2
    StandbyAsyncGroupName = rep_sa
    DefaultMode = STANDBY_ASYNC_TARGET
)

Mount mount_res (
    MountPoint = "/3dcmount"
    BlockDevice = "/dev/vx/dsk/3dctestdg/3dctestvol"
    FSType = vxfs
    FsckOpt = "-y"
)

app_res requires mount_res
dg_res requires htc3dc_res
mount_res requires dg_res

// resource dependency tree
//
// group htc3dcsg
// {
//   FileOnOff app_res
//   {
//     Mount mount_res
//     {
//       DiskGroup dg_res
//       {

```



```
//          HTC3DCASYNC htc3dc_res
//          }
//          }
//      }
// }
```

Before you configure the agent for Hitachi/HP XP 3 Data Center Asynchronous 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.
 See [“Typical Hitachi/HP XP 3 Data Center Asynchronous Replication setup in a VCS cluster”](#) on page 8.
- Make sure that the cluster has an effective heartbeat mechanism in place.
 See [“About cluster heartbeats”](#) on page 33.

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.

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 Aysnc 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 Asynchronous Replication

You can adapt most clustered applications to a disaster recovery environment by:

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

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

See [Figure 3-1](#) on page 23.

Configuring the agent to compute RPO

In a global cluster environment, the agent for Hitachi/HP XP 3 Data Center Asynchronous 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 HTC3DCASync resource and the VxVM disk group resource are online at the production site.
- The disk group resource is dependent on the HTC3DCASync 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 HTC3DCASYNC_resource_name ComputeDRSLA 1 -sys
system_name.
```

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

```
hares -action HTC3DCASYNC_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 HTC3DCASYNC_resource_name ComputeDRSLA 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:

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