

Cluster Server Agent for IBM XIV Mirror Installation and Configuration Guide

AIX, Linux, Solaris

7.0

Veritas InfoScale™ Availability Agents

Last updated: 2018-07-06

Legal Notice

Copyright © 2018 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 IBM XIV Mirror	6
	About the agent for IBM XIV Mirror	6
	Supported software	6
	Supported hardware for IBM XIV Mirror	6
	Typical IBM XIV Mirror setup in a VCS cluster	7
	IBM XIV Mirror agent functions	7
	About the agent's online function	9
Chapter 2	Installing and removing the agent for IBM XIV Mirror	11
	Before you install the agent for IBM XIV Mirror	11
	Installing the agent for IBM XIV Mirror	11
	Upgrading the agent for IBM XIV Mirror	12
	Removing the agent for IBM XIV Mirror	14
Chapter 3	Configuring the agent for IBM XIV Mirror	15
	Configuration concepts for the IBM XIV Mirror agent	15
	Resource type definition for the IBM XIV Mirror agent	15
	Attribute definitions for the XIV Mirror agent	16
	Sample configuration for the IBM XIV Mirror agent	18
	Before you configure the agent for IBM XIV Mirror	20
	About cluster heartbeats	20
	About preventing split-brain	21
	Configuring the agent for IBM XIV Mirror	21
	Configuring the agent manually in a replicated data cluster	22
	Configuring the agent manually in a global cluster	22
	Encrypting passwords for the local and remote XIV arrays	22
	Configuring the agent to compute RPO	23

Chapter 4	Managing and testing clustering support for IBM XIV Mirror	25
	How VCS recovers from various disasters in an HA/DR setup with IBM	
	XIV Mirror	25
	Failure scenarios in global clusters	26
	Failure scenarios in replicated data clusters	28
	Testing the global service group migration	31
	Testing disaster recovery after host failure	32
	Testing disaster recovery after site failure	33
	Performing failback after a node failure or an application failure	35
	Performing failback after a site failure	35

Introducing the agent for IBM XIV Mirror

This chapter includes the following topics:

- [About the agent for IBM XIV Mirror](#)
- [Supported software](#)
- [Supported hardware for IBM XIV Mirror](#)
- [Typical IBM XIV Mirror setup in a VCS cluster](#)
- [IBM XIV Mirror agent functions](#)

About the agent for IBM XIV Mirror

The Cluster Server (VCS) agent for IBM XIV Mirror manages the roles of the IBM XIV Mirror devices that have been configured for synchronous and asynchronous replication.

Supported software

For information on the software versions that the agent for IBM XIV Mirror supports, see the Veritas Services and Operations Readiness Tools (SORT) site: <https://sort.veritas.com/agents>.

Supported hardware for IBM XIV Mirror

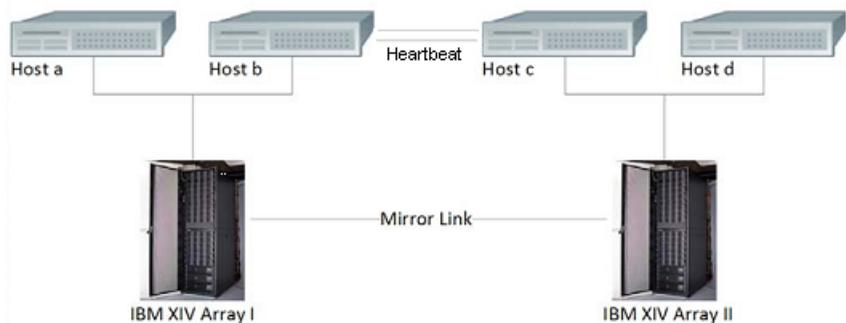
IBM XIV array version: 10.2.2.a and later

IBM XIV Remote Mirror replications must be planned and established between XIV volumes or consistency groups.

Typical IBM XIV Mirror setup in a VCS cluster

A typical IBM XIV Mirror setup in a VCS cluster includes at least two hosts per IBM XIV array, a heartbeat connection between hosts of different arrays, and a mirror link between the arrays.

Figure 1-1 Typical IBM XIV Mirror setup in a VCS cluster



This configuration maintains the following links:

- Host a and host b are attached via fibre channel or iSCSI to XIV Array I.
- Host c and host d are attached via fibre channel or iSCSI to XIV Array II.
- Volumes or a consistency group from Array I are mirrored against volumes or a consistency group from Array II.
- 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 20.
- In a replicated data cluster environment, all hosts are part of the same cluster. You must connect them with the dual and dedicated networks that support LLT.
- In a global cluster environment, you must attach all hosts in a cluster to the same XIV array.

IBM XIV Mirror agent functions

The agent performs the following functions:

online	<p>If the state of all local devices is read-write enabled, the agent creates a lock file on the local host to indicate that the resource is online.</p> <p>In other cases, when the consistency group or all local volumes or are Slave, the online function takes various actions depending on the replication link state and replication active state. See “About the agent’s online function” on page 9.</p>
offline	Removes the lock file from the local host.
monitor	<p>Verifies that the lock file exists.</p> <p>If the lock file exists, the monitor function reports the status of the resource as online.</p> <p>If the lock file does not exist, the monitor function reports the status of the resource as offline.</p>
clean	Removes the lock file from the local host.
open	Removes the lock file from the host where the function is called. This action prevents a potential concurrency violation if the service group fails over to another node. Note that the agent does not remove the lock file if the agent was started after a <code>hastop -force</code> command.
info	Modifies or adds the resource information of the XIVMirror resource for showing consolidated information about local volumes or the consistency group. This includes information about the role, synchronization type, link state, mirror state, and sync state.

action/PreSwitch	<p>Ensures that the remote site within a global cluster can come online during a planned failover.</p> <p>During a planned failover, the VCS engine at the remote site invokes the PreSwitch action on all the resources of the remote site by using the <code>hagrp -switch</code> command. For this, the PreSwitch attribute must be set to 1.</p> <p>The option <code>-nopre</code> indicates that the VCS engine must switch the service group regardless of the value of the PreSwitch service group attribute. The operation exits with an error if the replication link is down, or the replication is inactive, or the data at secondary site is not up-to-date.</p> <p>For synchronous replication, the agent checks if the <code>sync_state</code> of the replicated data is Consistent.</p> <p>For asynchronous replication, the agent checks if the <code>sync_state</code> of the replicated data is RPO OK.</p> <p>The operation exits if the state is not as expected. If the PreSwitch action fails to run successfully, failover does not take place. This minimizes application downtime and data loss.</p>
action/GetCurrentRPO	<p>Fetches the current point in time RPO. The agent performs this action function on the disaster recovery (DR) system where the <code>ComputeDRSLA</code> attribute is set to 1. 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

The online function attempts to make the specified XIV devices read-writable.

When the consistency group or all the local volumes are Master, the agent creates the lock file and exits.

When the consistency group or all the local volumes are Slave, the mirror link state and the replication active state determine the action the agent takes:

- If the mirror link state is connected, the replication state is active, the agent takes actions according to the mode of replication.

- If the mode of replication is synchronous and mirror sync state is Consistent, the agent executes the switch role operation sequentially for all the volumes or the consistency group.
- If the mode of replication is asynchronous and the mirror sync state is RPO OK, the agent takes the following steps:
 - Waits until the OnlineTimeout value is about to expire and checks if the consistency group or any of the volumes have sync jobs.
 - If the consistency group or any of the volumes have sync jobs, the agent logs an error message and does not initialize the switch role operation for all the mirrors. This ensures that the XIV volumes or consistency group do not end in mixed replication roles due to potential failure of the switch role command because a mirror has a sync job.
 - If there is no consistency group or volume with sync jobs, the agent executes the switch role operation sequentially for all the volumes or the consistency group.
 If, for any of the volumes, this operation fails with errors MIRROR_HAS_SYNC_JOB or MIRROR_RETRY_OPERATION, the agent retries the switch role operation for those volumes after some time, before the OnlineTimeout value is about to expire.
- If the mirror link state is not connected, and the AutoTakeover attribute is 0, the agent does not take any action.
- If the mirror link state is not connected, and the AutoTakeover attribute is 1, the agent executes the change role command.
- If the mirror replication state is inactive, and the SplitTakeover attribute is 0, the agent does not take any action.
- If the mirror replication state is inactive, and the SplitTakeover attribute is 1, the agent executes the change role command. After successful execution of the switch role or change role command, the agent creates the lock file.

If the local volumes have different roles, replication states, replication modes (mix of synchronous and asynchronous replication), and link states, the agent logs an error and does not take any action.

Installing and removing the agent for IBM XIV Mirror

This chapter includes the following topics:

- [Before you install the agent for IBM XIV Mirror](#)
- [Installing the agent for IBM XIV Mirror](#)
- [Upgrading the agent for IBM XIV Mirror](#)
- [Removing the agent for IBM XIV Mirror](#)

Before you install the agent for IBM XIV Mirror

Before you install the VCS agent for IBM XIV Mirror, 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*.

Install IBM XCLI on each of the cluster nodes on which you plan to install the IBM XIV agent.

Installing the agent for IBM XIV Mirror

You must install the IBM XIV Mirror 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 Veritas Services and Operations Readiness Tools (SORT) site: <https://sort.veritas.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.

```
AIX      cdl/aix/vcs/replication/xivmirror_agent/  
        agent_version/pkgs/
```

```
Linux   cdl/linux/generic/vcs/replication/xivmirror_agent/  
        agent_version/rpms/
```

```
Solaris cdl/solaris/dist_arch/vcs/replication/xivmirror_agent/  
        agent_version/pkgs/
```

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

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

```
AIX      # installp -ac -d VRTSvcsxivm.rte.bff VRTSvcsxivm.rte
```

```
Linux    # rpm -ihv \  
        VRTSvcsxivm-AgentVersion-Linux_GENERIC.noarch.rpm
```

```
Solaris  # pkgadd -d . VRTSvcsxivm
```

Upgrading the agent for IBM XIV Mirror

You must upgrade the agent on each node in the cluster.

To upgrade the agent software

- 1 Specify the location of the VCS binary in the PATH variable.

```
# export PATH=PATH:/opt/VRTSvcs/bin
```

- 2 Save the VCS configuration and stop the VCS engine.

```
# haconf -dump -makero
```

```
# haagent -stop XIVMirror -force -sys system
```

- 3 Verify the status of the agent.

```
# haagent -display XIVMirror
```

- 4 Remove the previous version of the agent from the node.

See [“Removing the agent for IBM XIV Mirror”](#) on page 14.

- 5 Install the latest version of the agent.

See [“Installing the agent for IBM XIV Mirror”](#) on page 11.

- 6 Start the agent.

```
# haagent -start XIVMirror
```

- 7 Verify the status of the agent.

```
# haagent -display XIVMirror
```

- 8 If a consistency group is already available for the mirror pairs that the resource manages, run the following commands to add the consistency group name to the GroupName attribute value.

```
# haconf -makerw
```

```
# hares -modify res_name GroupName CG_name
```

- 9 To continue monitoring the resource using the LocalVolNames attribute value [Local XIV volumes], run the following commands.

```
# haconf -makerw
```

```
# hares -modify res_name IsConsistencyGroup 0
```

Note: Before you bring previously-configured XIV Mirror resources online, consider making the following changes to the attribute values:

If a consistency group is already available for the mirror pairs that the resource manages, and if you want to leverage the agent's ability to support a consistency group, ensure that you add the consistency group name to the GroupName attribute value.

If you want the agent to continue to manage the mirror pairs as specified in the LocalVolNames attribute value, make sure that you set the value of the IsConsistencyGroup attribute to 0.

Veritas recommends that you use consistency groups for disaster recovery operations through the XIVMirror agent.

Removing the agent for IBM XIV Mirror

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

You must remove the XIV Mirror agent from each node in the cluster.

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

```
AIX          # installp -u VRTSvcsxivm.rte
```

```
Linux       # rpm -e VRTSvcsxivm
```

```
Solaris    # pkgrm VRTSvcsxivm
```

Note: To uninstall the agent IPS package on a Solaris 11 system:

```
# pkg uninstall VRTSvcsxivm
```

Configuring the agent for IBM XIV Mirror

This chapter includes the following topics:

- [Configuration concepts for the IBM XIV Mirror agent](#)
- [Before you configure the agent for IBM XIV Mirror](#)
- [Configuring the agent for IBM XIV Mirror](#)

Configuration concepts for the IBM XIV Mirror agent

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

Resource type definition for the IBM XIV Mirror agent

The resource type definition for the XIV Mirror agent is as follows:

```
type XIVMirror (
    static keylist RegList = { ComputeDRSLA }
    static keylist SupportedActions = { PreSwitch, ReportRPOData,
                                        StartWriter, GetCurrentRPO,
                                        StartRPOComputation,
                                        StopRPOComputation }

    static int OpenTimeout = 180
    static int RestartLimit = 1
    static str ArgList[] = {XCliPath, UserLocal, PasswordLocal,
                            UserRemote, PasswordRemote, LocalXIVIPs,
                            RemoteXIVIPs, LocalVolNames, AutoTakeover,
```

```

SplitTakeover, IsConsistencyGroup,
GroupName, ComputeDRSLA}

str XcliPath
str UserLocal
str PasswordLocal
str UserRemote
str PasswordRemote
str LocalXIVIPs[]
str RemoteXIVIPs[]
str LocalVolNames[]
int AutoTakeover = 1
int SplitTakeover = 0
int IsConsistencyGroup = 1
str GroupName
temp str VCSResLock
int ComputeDRSLA
temp boolean Tagging = 0
)

```

Attribute definitions for the XIV Mirror agent

The attributes for the XIV Mirror agent are as follows:

XcliPath	Specifies the XIV command line interface name with its full path.
LocalXIVIPs	List of local XIV IP addresses to connect.
RemoteXIVIPs	List of remote XIV IP addresses to connect.
LocalVolNames	List of local XIV volume names that the agent manages.
AutoTakeover	Indicates whether the agent should enable read/write access to the local devices in the replication relationship when the replication link is broken. If it is set to 0, the agent does not enable read/write access when the replication link is broken.
UserLocal	Specifies the user for the local XIV array.
PasswordLocal	Encrypted password for the local XIV array. Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the <code>vcscrypt</code> utility. See “Encrypting passwords for the local and remote XIV arrays” on page 22.
UserRemote	Specifies the user for the remote XIV array.

PasswordRemote	<p>Encrypted password for the remote XIV array.</p> <p>Encrypt passwords only when entering them using the command-line. Passwords must be encrypted using the <code>vcseencrypt</code> utility. See “Encrypting passwords for the local and remote XIV arrays” on page 22.</p>
SplitTakeover	<p>Indicates whether the agent should enable read/write access to the local devices in the replication relationship when the replication is in an inactive state. If it is set to 0, the agent does not enable read/write access when the replication is in an inactive state.</p>
IsConsistencyGroup	<p>Specifies if the agent must manage a consistency group or a list of mirror pairs.</p> <p>When the value of this attribute is set to 0, it means that the agent must manage mirror pairs. When the value of this attribute is set to 1, it means that the agent must manage a consistency group.</p> <p>Note: If you configure both, the <code>LocalVolNames</code> attribute and the <code>GroupName</code> attribute, the agent refers to the value of the <code>IsConsistencyGroup</code> attribute to determine if the agent must manage mirror pairs or a consistency group.</p> <p>Default Value: 1</p>
GroupName	<p>The name of the consistency group that the agent must manage.</p> <p>A resource can manage only one consistency group.</p>
ComputeDRSLA	<p>Used to enable or disable Recovery Point Objective (RPO) computation.</p> <p>Set this attribute on any one node in the disaster recovery (DR) cluster. Setting this attribute to 1 starts the RPO computation process.</p> <p>Ensure that you reset this attribute to 0 after you use the <code>GetCurrentRPO</code> action function to check the RPO.</p> <p>Type-dimension : integer-scalar</p> <p>Default value : 0</p>
Tagging	<p>This internal attribute is used for maintaining the process of computing RPO.</p>

Consider the following points:

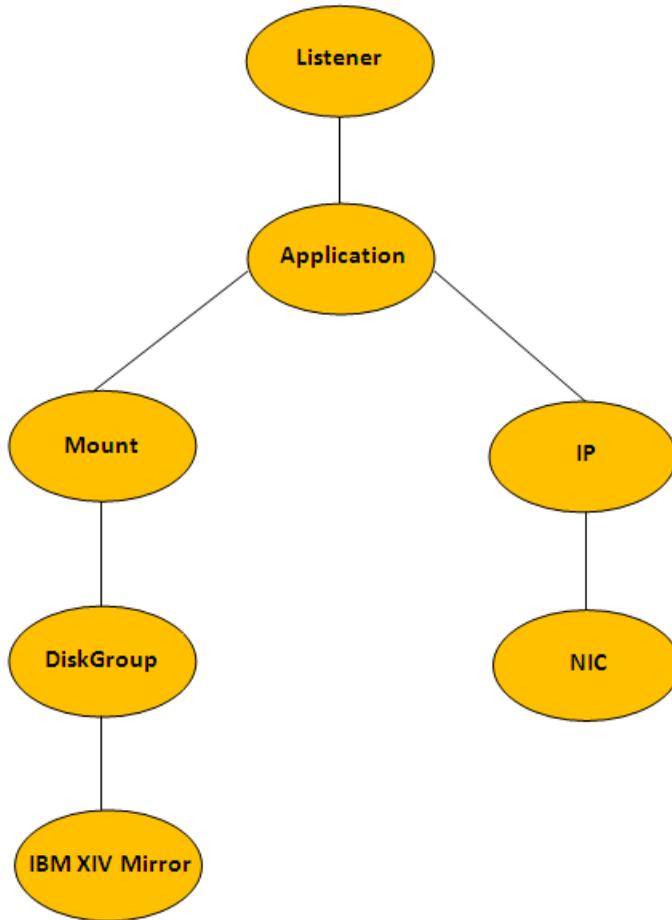
- If you plan to use the environment variables—`XIV_XCLIUSER` and `XIV_XCLIPASSWORD`, these variables must be set before HAD or VCS starts.

- If you are using XCLI version 3.1.1.x or later, Veritas recommends using the `xcli --protect` command to encrypt the credentials of the local and remote arrays.
- The `UserLocal` and `PasswordLocal` attributes are optional when any one of the following conditions are met:
 - The environment variables—`XIV_XCLIUSER` and `XIV_XCLIPASSWORD`—on the host, point to the username and password of the local XIV array.
 - The local array credentials are encrypted by using the `xcli --protect` command.
- The `UserRemote` and `PasswordRemote` attributes are optional when any one of the following conditions are met:
 - The environment variables—`XIV_XCLIUSER` and `XIV_XCLIPASSWORD`—point to the username and password of the remote XIV array.
 - The remote array credentials are encrypted by using the `xcli --protect` command.

Sample configuration for the IBM XIV Mirror agent

[Figure 3-1](#) shows the dependency graph for a VCS service group with a resource of type `XIVMirror`.

Figure 3-1 Sample configuration for the XIV Mirror agent



You can configure a resource of type XIVMirror in the main.cf file.

A sample `main.cf` file when the agent is configured to manage volume mirror pairs is as follows:

```

XIVMirror xiv (
    Critical = 0
    XcliPath = "/root/IBM_XIV_Storage_Management_GUI/xcli"
    UserLocal = LocalUser
    PasswordLocal = LocalPassword
    UserRemote = RemoteUser
    PasswordRemote = RemotePassword
    LocalXIVIPs = { "XXX.XXX.XXX.XXX" }
  )

```

```
RemoteXIVIPs = { "YYY.YYY.YYY.YYY" }
LocalVolNames = { Volume_1, Volume_2 }
SplitTakeover = 1
IsConsistencyGroup = 0
)
```

A sample `main.cf` file when the agent is configured to manage a consistency group is as follows:

```
XIVMirror xiv (
    Critical = 0
    XCliPath = "/root/IBM_XIV_Storage_Management_GUI/xcli"
    UserLocal = LocalUser
    PasswordLocal = LocalPassword
    UserRemote = RemoteUser
    PasswordRemote = RemotePassword
    LocalXIVIPs = { "XXX.XXX.XXX.XXX" }
    RemoteXIVIPs = { "YYY.YYY.YYY.YYY" }
    SplitTakeover = 1
    GroupName = VCS_CG_Sync
    IsConsistencyGroup = 1
)
```

Before you configure the agent for IBM XIV Mirror

Before you configure the agent, review the following information:

- Verify that you have installed the agent on all systems in the cluster.
- Verify that IBM XCLI is installed on each cluster node on which you have installed the agent.
- Verify the hardware setup for the agent.
See [“Typical IBM XIV Mirror setup in a VCS cluster”](#) on page 7.
- Make sure that the cluster has an effective heartbeat mechanism in place.
See [“About cluster heartbeats”](#) on page 20.

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.

About preventing split-brain

Split-brain occurs when all heartbeat links between the primary and secondary hosts are cut. In this situation, each side mistakenly assumes that the other side is down. You can minimize the effects of split-brain by ensuring that the cluster heartbeat links pass through a similar physical infrastructure as the replication links. When you ensure that both pass through the same infrastructure, if one breaks, so does the other.

Sometimes you cannot place the heartbeats alongside the replication links. In this situation, a possibility exists that the cluster heartbeats are disabled, but the replication link is not. A failover transitions the original primary to secondary and secondary to primary. In this case, the application faults because its underlying volumes become write-disabled, causing the service group to fault. VCS tries to fail it over to another host, causing the same consequence in the reverse direction. This phenomenon continues until the group comes online on the final node. You can avoid this situation by setting up your infrastructure such that loss of heartbeat links also means the loss of replication links.

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

Configuring the agent for IBM XIV Mirror

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

- Converting their devices to XIV Mirror devices
- Synchronizing the devices
- Adding the IBM XIV Mirror agent to the service group

Configure IBM XIV mirrored volumes as resources of type XIVMirror.

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

See [“Sample configuration for the IBM XIV Mirror agent”](#) on page 18.

Note: You must not change the replication state of devices from primary to secondary and from secondary to primary, outside of a VCS setup. The agent for IBM XIV Mirror fails to detect a change in the replication state if the role reversal is done externally and RoleMonitor is disabled.

Configuring the agent manually in a replicated data cluster

Configuring the agent manually in a replicated data cluster involves the following tasks:

To configure the agent in a replicated data cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 In each service group that uses replicated data, add a resource of type XIVMirror at the bottom of the service group.
- 3 Configure the attributes of the XIVMirror resource. Note that some attributes must be localized to reflect values for the hosts that are attached to different arrays.
- 4 Set the SystemZones attribute for the service group to reflect which hosts are attached to the same array.

Configuring the agent manually in a global cluster

Configuring the agent manually in a global cluster involves the following tasks:

To configure the agent in a global cluster

- 1 Start Cluster Manager and log on to the cluster.
- 2 Add a resource of type XIVMirror at the bottom of the service group.
- 3 Configure the attributes of the XIVMirror resource.
- 4 If the service group is not configured as a global service group, configure the service group using the Global Group Configuration Wizard.
- 5 Repeat step 2 through step 4 for each service group in each cluster that uses replicated data.
- 6 The configuration must be identical on all cluster nodes, both primary and disaster recovery.

Encrypting passwords for the local and remote XIV arrays

If you choose to use XCLI to specify the user names and passwords for the local and remote XIV arrays, the passwords must be specified and stored in the VCS configuration in an encrypted format. You must use the VCS-provided `vcseencrypt` utility to encrypt the passwords for the local and remote XIV arrays, and you must then specify these encrypted passwords when you configure the PasswordLocal and PasswordRemote attributes, respectively.

For more information about the `vcseencrypt` utility, see the `vcseencrypt(1M)` manual page.

The `vcseencrypt` utility also allows you to encrypt the agent passwords using a security key. The security key supports Advanced Encryption Standard (AES) encryption which creates a more secure password for the agent. For details, see the *Cluster Server Administrator's Guide*.

To encrypt passwords without using security keys

- 1 Specify the location of the VCS binary in the PATH variable.

```
# export PATH=PATH:/opt/VRTSvcs/bin
```

- Type the following command:

```
# vcseencrypt -agent
```

- Enter the password and confirm it by entering it again. Press Enter.

```
Enter New Password:
```

```
Enter Again:
```

- 2 Run the `vcseencrypt` utility.

The `vcseencrypt` utility encrypts the password and displays the encrypted password.

- 3 Enter this encrypted password as the value for the attribute.
- 4 Copy the encrypted password for future reference.

Configuring the agent to compute RPO

In a global cluster environment, the agent for IBM XIV Mirror 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 XIVMirror resource and the VxVM disk group resource are online at the production site.
- The disk group resource is dependent on the XIVMirror 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 XIVMirror_resource_name ComputeDRSLA 1 -sys  
system_name.
```

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

```
hares -action XIVMirror_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 XIVMirror_resource_name ComputeDRSLA 0 -sys  
system_name
```

Managing and testing clustering support for IBM XIV Mirror

This chapter includes the following topics:

- [How VCS recovers from various disasters in an HA/DR setup with IBM XIV Mirror](#)
- [Testing the global service group migration](#)
- [Testing disaster recovery after host failure](#)
- [Testing disaster recovery after site failure](#)
- [Performing failback after a node failure or an application failure](#)
- [Performing failback after a site failure](#)

How VCS recovers from various disasters in an HA/DR setup with IBM XIV Mirror

This topic lists various failure scenarios and describes how VCS responds to the failures in the following DR cluster configurations.

Global clusters

When a site-wide global service group or system fault occurs, VCS failover behavior depends on the value of the ClusterFailOverPolicy attribute for the faulted global service group. The VCS agent for IBM XIV Mirror ensures safe and exclusive access to the configured IBM XIV Mirror devices.

See [“Failure scenarios in global clusters”](#) on page 26.

Replicated data clusters

When service group faults or system faults occur, the VCS failover behavior depends on the value of the AutoFailOver attribute of the faulted service group. The VCS agent for IBM XIV Mirror ensures safe and exclusive access to the configured IBM XIV Mirror devices.

See [“Failure scenarios in replicated data clusters”](#) on page 28.

Refer to the *Cluster Server Administrator's Guide* for more information on the DR configurations and the global service group attributes.

Failure scenarios in global clusters

The following table lists the failure scenarios in a global cluster configuration and describes the behavior of VCS and the agent in response to the failure.

Table 4-1 Failure scenarios in a global cluster configuration with the VCS agent for IBM XIV Mirror

Failure	Description and VCS response
Application failure	<p>Application cannot start successfully on any hosts at the primary site.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Causes global service group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ Auto or Connected—VCS automatically brings the faulted global group online at the secondary site. ■ Manual—No action. You must bring the global group online at the secondary site. <p>Agent response:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ If replication is suspended and the value of the SplitTakeover attribute is set to 0, the agent takes no action. ■ If the write mode is asynchronous, sync state is RPO_OK, and SplitTakeover is set to 1, the agent issues failover. ■ Promotes the volumes at the remote site to Master. <p>See “Performing failback after a node failure or an application failure” on page 35.</p>

Table 4-1 Failure scenarios in a global cluster configuration with the VCS agent for IBM XIV Mirror (*continued*)

Failure	Description and VCS response
Host failure	<p>All hosts at the primary site fail.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Displays an alert to indicate the primary cluster fault. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ Auto—VCS automatically brings the faulted global group online at the secondary site. ■ Manual or Connected—No action. You must bring the global group online at the secondary site. <p>The agent does the following:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ In case the replication is suspended and the value of the SplitTakeover attribute is set to 0, no action is taken by agent. ■ In case the write mode is asynchronous, sync state is RPO_OK and SplitTakeover is set to 1, the agent will issue failover. ■ Promotes the volumes at the remote site to Master. <p>See “Performing failback after a node failure or an application failure” on page 35.</p>
Site failure	<p>All hosts and the storage at the primary site fail.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Displays an alert to indicate the cluster fault. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ Auto—VCS automatically brings the faulted global group online at the secondary site. ■ Manual or Connected—No action. You must bring the global group online at the secondary site. <p>Agent response: The agent does the following based on the value of the AutoTakeover attribute of the XIVMirror resource:</p> <ul style="list-style-type: none"> ■ 1—The agent issues the <code>failover</code> command to promote the volumes at the remote site to Master. ■ 0—No action is taken by the agent. The XIVMirror resource is faulted. <p>See “Performing failback after a site failure” on page 35.</p>

Table 4-1 Failure scenarios in a global cluster configuration with the VCS agent for IBM XIV Mirror (*continued*)

Failure	Description and VCS response
Replication link failure	<p>Replication link between the arrays at the two sites fails.</p> <p>VCS response: No action.</p> <p>Agent response: The agent does the following based on the value of the AutoTakeover attribute of the XIVMirror resource:</p> <ul style="list-style-type: none"> ■ 1—The agent issues the <code>failover</code> command to promote the volumes at the remote site to Master. ■ 0—No action is taken by the agent. The XIVMirror resource is faulted.
Network failure	<p>The network connectivity and the replication link between the sites fail.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ VCS at each site concludes that the remote cluster has faulted. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ Manual or Connected—No action. You must confirm the cause of the network failure from the cluster administrator at the remote site and fix the issue. ■ Auto—VCS brings the global group online at the secondary site which may lead to a site-wide split brain. This causes data divergence between the devices on the primary and the secondary arrays. <p>When the network (WAC and replication) connectivity is restored, you must manually resync the data.</p> <p>Note: Veritas recommends that the value of the ClusterFailOverPolicy attribute is set to Manual for all global groups to prevent unintended failovers due to transient network failures.</p>
Storage failure	<p>The array at the primary site fails.</p> <p>VCS response at the secondary site:</p> <ul style="list-style-type: none"> ■ Causes the global service group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the ClusterFailOverPolicy global service group attribute: <ul style="list-style-type: none"> ■ Auto or Connected—VCS automatically brings the faulted global service group online at the secondary site. ■ Manual—No action. You must bring the global group online at the secondary site.

Failure scenarios in replicated data clusters

The following table lists the failure scenarios in a replicated data cluster configuration, and describes the behavior of VCS and the agent in response to the failure.

Table 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for IBM XIV Mirror

Failure	Description and VCS response
Application failure	<p>Application cannot start successfully on any hosts at the primary site.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at the secondary site. ■ 2—You must bring the service group online at the secondary site. <p>The agent does the following:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ In case the replication is suspended and the value of the SplitTakeover attribute is set to 0, no action is taken by agent. ■ In case the write mode is asynchronous, sync state is RPO_OK and SplitTakeover is set to 1, the agent will issue failover. ■ Promotes the volumes at the remote site to Master. <p>See “Performing failback after a node failure or an application failure” on page 35.</p>
Host failure	<p>All hosts at the primary site fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at the secondary site. ■ 2—You must bring the service group online at the secondary site. <p>The agent does the following:</p> <ul style="list-style-type: none"> ■ Write enables the devices at the secondary site. ■ In case the replication is suspended and the value of the SplitTakeover attribute is set to 0, no action is taken by agent. ■ In case the write mode is asynchronous, sync state is RPO_OK and SplitTakeover is set to 1, the agent will issue failover. ■ Promotes the volumes at the remote site to Master. <p>See “Performing failback after a node failure or an application failure” on page 35.</p>

Table 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for IBM XIV Mirror (*continued*)

Failure	Description and VCS response
Site failure	<p>All hosts and the storage at the primary site fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at the secondary site. ■ 2—You must bring the service group online at the secondary site. <p>Agent response: The agent does the following based on the value of the AutoTakeover attribute of the XIVMirror resource:</p> <ul style="list-style-type: none"> ■ 1—The agent issues the <code>failover</code> command to promote the volumes at the remote site to Master. ■ 0—No action is taken by the agent. The XIVMirror resource is faulted. <p>See “Performing failback after a site failure” on page 35.</p>
Replication link failure	<p>Replication link between the arrays at the two sites fails.</p> <p>VCS response: No action.</p> <p>Agent response: The agent does the following based on the value of the AutoTakeover attribute of the XIVMirror resource:</p> <ul style="list-style-type: none"> ■ 1—The agent issues the <code>failover</code> command to promote the volumes at the remote site to Master. ■ 0—No action is taken by the agent. The XIVMirror resource is faulted.

Table 4-2 Failure scenarios in a replicated data cluster configuration with VCS agent for IBM XIV Mirror (*continued*)

Failure	Description and VCS response
Network failure	<p>The LLT and the replication links between the sites fail.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ VCS at each site concludes that the nodes at the other site have faulted. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 2—No action. You must confirm the cause of the network failure from the cluster administrator at the remote site and fix the issue. ■ 1—VCS brings the service group online at the secondary site which leads to a cluster-wide split brain. This causes data divergence between the devices on the arrays at the two sites. <p>When the network (LLT and replication) connectivity is restored, VCS takes all the service groups offline on one of the sites and restarts itself. This action eliminates concurrency violation where in the same group is online at both the sites.</p> <p>Note: Veritas recommends that the value of the AutoFailOver attribute is set to 2 for all service groups to prevent unintended failovers due to transient network failures.</p> <p>Agent response: Similar to the site failure.</p>
Storage failure	<p>The array at the primary site fails.</p> <p>VCS response:</p> <ul style="list-style-type: none"> ■ Causes the service group at the primary site to fault and displays an alert to indicate the fault. ■ Does the following based on the AutoFailOver attribute for the faulted service group: <ul style="list-style-type: none"> ■ 1—VCS automatically brings the faulted service group online at the secondary site. ■ 2—You must bring the service group online at the secondary site.

Testing the global service group migration

After you configure the Cluster Server agent for IBM XIV Mirror, verify that the global service group can migrate to hosts across the sites. Depending on your DR configuration, perform one of the following procedures.

To test the global service group migration in global cluster setup

- 1 Fail over the global service group from the primary site to the secondary site.

Perform the following steps:

- Switch the global service group from the primary site to any node in the secondary site.

```
hagrp -switch global_group -any -clus cluster_name
```

VCS brings the global service group online on a node at the secondary site.

- 2 Fail back the global service group from the secondary site to the primary site.

Perform the following steps:

- Switch the global service group from the secondary site to the primary site.

```
hagrp -switch global_group -any -clus cluster_name
```

VCS brings the global service group online at the primary site.

To test service group migration in replicated data cluster setup

- 1 Fail over the service group from the primary site to the secondary site.

Perform the following steps:

- Switch the service group from the primary site to any node in the secondary site.

```
hagrp -switch service_group -to sys_name
```

VCS brings the service group online on a node at the secondary site.

- 2 Fail back the service group from the secondary site to the primary site.

Perform the following steps:

- Switch the service group from the secondary site to any node in the primary site.

```
hagrp -switch service_group -to sys_name
```

VCS brings the service group online on a node at the primary site.

Testing disaster recovery after host failure

Review the details on host failure and how VCS and the Cluster Server agent for IBM XIV Mirror behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 26.

See [“Failure scenarios in replicated data clusters”](#) on page 28.

Perform the procedure that is applicable to your DR configuration to test how VCS recovers after all hosts at the primary site fail.

To test disaster recovery for host failure in global cluster setup

- 1 Halt the hosts at the primary site.

The value of the ClusterFailOverPolicy attribute for the faulted global group determines the VCS failover behavior.

- Auto—VCS brings the faulted global service group online at the secondary site.
- Manual or Connected—You must bring the global service group online at the secondary site.

On a node in the secondary site, run the following command:

```
hagrp -online -force global_group -any
```

- 2 Verify that the global service group is online at the secondary site.

```
hagrp -state global_group
```

To test disaster recovery for host failure in replicated data cluster setup

- 1 Halt the hosts at the primary site.

The value of the AutoFailOver attribute for the faulted service group determines the VCS failover behavior.

- 1—VCS brings the faulted service group online at the secondary site.
- 2—You must bring the service group online at the secondary site.

On a node in the secondary site, run the following command:

```
hagrp -online service_group -to sys_name
```

- 2 Verify that the service group is online at the secondary site.

```
hagrp -state global_group
```

Testing disaster recovery after site failure

Review the details on site failure and how VCS and the Cluster Server agent for IBM XIV Mirror behave in response to the failure.

See [“Failure scenarios in global clusters”](#) on page 26.

See [“Failure scenarios in replicated data clusters”](#) on page 28.

Depending on the DR configuration, perform one of the following procedures to test the disaster recovery in the event of site failure.

To test disaster recovery for site failure in global cluster setup

- 1 Halt all nodes and the arrays at the primary site.

If you cannot halt the array at the primary site, then disable the replication link between the two arrays.

The value of the ClusterFailOverPolicy attribute for the faulted global group determines the failover behavior of VCS.

- Auto—VCS brings the faulted global group online at the secondary site.
- Manual or Connected—You must bring the global group online at the secondary site.

On a node in the secondary site, run the following command:

```
hagrp -online -force global_group -any
```

- 2 Verify that the global service group is online at the secondary site.

```
hagrp -state global_group
```

To test disaster recovery for site failure in replicated data cluster setup

- 1 Halt all hosts and the arrays at the primary site.

If you cannot halt the array at the primary site, then disable the replication link between the two arrays.

The value of the AutoFailOver attribute for the faulted global service group determines the VCS failover behavior.

- 1—VCS brings the faulted global service group online at the secondary site.
- 2—You must bring the global service group online at the secondary site.

On a node in the secondary site, run the following command:

```
hagrp -online service_group -sys sys_name
```

- 2 Verify that the global service group is online at the secondary site.

```
hagrp -state global_group
```

Performing failback after a node failure or an application failure

Review the details on node failure and application failure and how VCS and the agent for IBM XIV Mirror behave in response to these failures.

See [“Failure scenarios in global clusters”](#) on page 26.

See [“Failure scenarios in replicated data clusters”](#) on page 28.

After the nodes at the primary site are restarted, you can perform a failback of the global service group to the primary site. Perform the procedure that applicable to your DR configuration.

To perform failback after a node failure or an application failure in global cluster

- ◆ Switch the global service group from the secondary site to any node in the primary site.

```
hagrp -switch global_group -any -clus cluster_name
```

VCS brings the global service group online at the primary site.

To perform failback after a host failure or an application failure in replicated data cluster

- ◆ Switch the global service group from the secondary site to any node in the primary site.

```
hagrp -switch service_group -to sys_name
```

VCS brings the global service group online on a node at the primary site.

Performing failback after a site failure

See [“Failure scenarios in global clusters”](#) on page 26.

See [“Failure scenarios in replicated data clusters”](#) on page 28.

To perform failback after a site failure in global cluster

- 1 Take the global service group offline at the secondary site. On a node at the secondary site, run the following command:

```
hagrp -offline global_group -any
```

- 2 Bring the global service group online at the primary site. On a node in the primary site, run the following command:

```
hagrp -online global_group -any
```

To perform failback after a site failure in replicated data cluster

- 1 Take the global service group offline at the secondary site. On a node in the secondary site, run the following command:

```
hagrp -offline service_group -sys sys_name
```

- 2 Bring the global service group online at the primary site. On a node in the primary site, run the following command:

```
hagrp -online service_group -sys sys_name
```