

HP StorageWorks Cluster Extension XP user guide

XP48
XP128
XP512
XP1024
XP10000
XP12000

product version: 2.06.00

seventh edition (October 2005)

part number T1609-96006

This guide explains how to use the HP StorageWorks Cluster Extension XP software.



© Copyright 2003-2005 Hewlett-Packard Development Company, L.P., all rights reserved

Hewlett-Packard Company makes no warranty of any kind with regard to this material, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. Hewlett-Packard shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this material.

This document contains proprietary information, which is protected by copyright. No part of this document may be photocopied, reproduced, or translated into another language without the prior written consent of Hewlett-Packard. The information contained in this document is subject to change without notice.

All product names mentioned herein may be trademarks of their respective companies.

Hewlett-Packard Company shall not be liable for technical or editorial errors or omissions contained herein. The information is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Hewlett-Packard Company products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

Printed in the U.S.A.

HP StorageWorks Cluster Extension XP User Guide

product version: 2.06.00

seventh edition (October 2005)

part number: T1609-96006

Contents

About this guide	9
Intended Audience	10
Supported arrays	10
Disk array firmware and software dependencies	11
Related documentation	12
Other resources	12
Getting help	13
HP technical support	13
HP storage web site	14
HP authorized reseller	14
Document conventions	15
Terminology conventions	16
Revision history	17
Warranty statement	19
1 Cluster Extension XP features	21
Integration into cluster software	22
Disaster tolerance through geographical dispersion	23
Disaster tolerance considerations	23
Disaster-tolerant architectures	24
Automated redirection and monitoring of mirrored Continuous Access XP pairs	31
Rolling disaster protection	32
What is a rolling disaster?	32
Recovering the disaster tolerant environment	33
Command line interface for easy integration	35
Graphical user interface	36
Quorum service (<i>Microsoft Cluster Service only</i>)	37

2	Cluster Extension XP processes and components	39
	Cluster Extension XP environments	40
	Cluster Extension XP execution	42
	Continuous Access XP and RAID Manager XP	43
	RAID Manager XP instances	45
	RAID Manager XP device groups	46
	Rolling disaster protection and Business Copy XP	47
	Integration with RAID Manager XP	48
	Integration with automatic recovery	49
	Integration with the pair/resync monitor	49
	Restoring server operation	50
	Example	50
	User configuration file	52
	Pair/resync monitor	53
	Force flag	55
	Pre-execution and post-execution programs	56
	Cluster Extension XP log facility	59
	Error return codes	60
	Quorum service for Microsoft Cluster Service	61
	Using the quorum service in a Microsoft Cluster Service environment	61
	Quorum processes	62
3	User configuration file and Cluster Extension XP objects	67
	The user configuration file	68
	File structure	69
	Specifying object values	70
	COMMON section objects	73
	APPLICATION section objects	75
	Basic configuration example	89
4	RAID Manager XP dependencies	91
	RAID Manager XP configuration	92
	RAID Manager XP configuration file	92
	Network considerations	95
	Command device considerations	95
	Start and stop the RAID Manager XP instances	96
	Takeover basic functionality test	97

5	Integration with HACMP	99
	Configuring resources	100
	Procedure for HACMP	101
	User configuration file for HACMP	103
	Bringing a resource group online	106
	Taking a resource group offline	107
	Deleting Cluster Extension XP	108
	Pair/resync monitor integration	109
	Timing considerations	112
	Failure behavior	114
	Restrictions for IBM HACMP with Cluster Extension XP	115
6	Integration with Microsoft Cluster Service	117
	Configuring the quorum service	118
	Configuring Cluster Extension XP resources	119
	Resource group and resource names	120
	Cluster Extension XP resource-specific parameters	120
	Setting non-Cluster Extension XP resource-specific parameters	121
	Adding a Cluster Extension XP resource	125
	Changing Cluster Extension XP resource properties	127
	Advanced properties	129
	Changing a resource name	130
	Adding dependencies on a Cluster Extension XP resource	131
	Bringing a Cluster Extension XP resource online	133
	Taking a Cluster Extension XP resource offline	135
	Deleting a Cluster Extension XP resource	136
	Pair/resync monitor integration	137
	Timing considerations for Microsoft Cluster Service	139
	Failure behavior with Microsoft Cluster Service	141
	Bouncing Resource Groups	141
	Unexpected offline conditions	141
	Restrictions for Microsoft Cluster Service with Cluster Extension XP	143
	Disaster-tolerant configuration example using a file share	144
	Administration	149
7	Integration with VCS	151
	Configuration of the Cluster Extension XP agent	152
	Configuring the Cluster Extension XP resource	156
	Cluster Extension resource types	156
	Resource type definition	157

Adding a Cluster Extension XP resource	158
Changing Cluster Extension XP attributes	160
Linking a Cluster Extension XP resource	162
Bringing a Cluster Extension XP resource online	163
Taking a Cluster Extension XP resource offline	165
Deleting a Cluster Extension XP resource	166
Pair/resync monitor integration	167
Timing considerations for VCS	168
Enable/disable service groups	170
Restrictions for VCS with Cluster Extension XP	171
Unexpected offline conditions	173

8 Integration with Serviceguard for Linux 175

Configuration of the Cluster Extension XP environment	176
Adding a Cluster Extension XP integration to an existing Serviceguard package	182
Starting a Serviceguard package with Cluster Extension XP	183
Halting a Serviceguard package with Cluster Extension XP	184
Deleting Cluster Extension XP from a Serviceguard package	185
Pair/resync monitor integration	187
Timing considerations for Serviceguard	190

9 Command line interface (CLI) 193

Configuring the CLI	195
Creating the Continuous Access environment and configuring RAID Manager	195
Timing considerations	195
Restrictions for customized Cluster Extension XP implementations	197
Creating and configuring the user configuration file	197
CLI Commands	198
clxrun	198
clxchkmon	200
clxqr	203

10 Troubleshooting 205

Start errors	207
Failover error handling	208

HACMP-specific error handling	209
Start errors	209
Failover errors	209
Microsoft Cluster Service-specific error handling	213
Solving quorum service problems	213
Resource start errors	215
Failover errors	216
VCS-specific error handling	218
Start errors	218
Failover errors	219
Serviceguard (SG-LX)-specific error handling	222
Start errors	222
Failover errors	222
Pair/resync monitor messages in syslog/errorlog/messages/Event Log	224
A Recovery procedures	227
XP disk pair states	228
Recovery sequence	230
Quorum service recovery (<i>Microsoft Cluster Service only</i>)	232
Single site failure recovery	232
Failure recovery if both sites have failed	234
Procedure for quorum service system cleanup	235
B Cluster Extension XP resource message catalog	237
C Cluster Extension XP quorum service message catalog	263
Quorum service Event Log messages	269
Glossary	271
Index	275

About this guide

This guide provides information about configuring and using HP StorageWorks Cluster Extension XP in an environment where clustered systems are connected to a disaster recovery array-based mirroring solution. Cluster Extension XP allows creation of dispersed multiplatform cluster configurations with the XP disk array. Cluster Extension XP enables cluster software to automatically failover applications where data is stored and continuously mirrored from a local to a remote disk array using HP StorageWorks Continuous Access XP. This guide describes the options you have to make your disaster tolerant environment as robust as possible to keep your data available at all times.

Because the XP family of disk arrays supports a broad range of operating systems and cluster software, Cluster Extension XP can be integrated with almost any disk array-supported cluster software. This guide provides you with the necessary information to create a disaster tolerant environment of two or more data centers utilizing the XP disk array and its Continuous Access XP remote mirroring feature.

Cluster Extension XP features

HP StorageWorks Cluster Extension XP enables monitoring of HP StorageWorks Continuous Access XP-mirrored disk pairs and allows access to the remote data copy if an application becomes unavailable on the local site. If the application service is restarted on the remote site, after the local (primary) application service has been shut down, Cluster Extension XP uses its internal database to check whether the current disk states allow automatic access to your data, based on consistency and concurrency considerations. Integrated in the cluster software or available as command line interface for your own integration, Cluster Extension XP ensures that the data can be accessed if necessary.

Cluster Extension XP software provides these key features:

- integration into cluster software
- disaster tolerance through geographical dispersion
- automated redirection and monitoring of mirrored Continuous Access XP pairs
- command line interface for easy integration

Integration into cluster software

Cluster Extension XP provides tight integration with the cluster software wherever possible. Cluster Extension XP is a resource of the clustered application service (like the disk or volume group) and must be managed as such. The architecture of Cluster Extension XP allows integration with many cluster software products:

- VERITAS Cluster Server (VCS)
- IBM HACMP
- Windows 2000 Advanced Server and Datacenter Server Cluster service
- Windows Server 2003 Enterprise Edition and Datacenter Edition
- Serviceguard for Linux (SG-LX)

For the current list of supported cluster software, contact your HP representative.

Cluster Extension XP processes and components

Cluster Extension XP is shipped in the appropriate format for each platform:

Platform	Implementation
VCS	agent
IBM HACMP	pre-event executable
Microsoft Cluster Service	resource DLL, quorum service, and external arbitrator
SG-LX	function call/executable

Customized solutions to failover application services must implement Cluster Extension XP through its command line interface prior to the disk activation procedure.

Cluster Extension XP environments

The ideal environment for a Cluster Extension XP configuration consists of at least four servers (two at each site) and separated redundant communications links for cluster heartbeats, client access and Continuous Access XP (Extension). All communications interfaces must be installed in pairs to serve as failover components, preventing single points of failure (SPOFs).

Recommendation Use load balancing and alternative pathing software for host-to-storage connections, such as HP StorageWorks Auto Path for IBM AIX or Secure Path for Linux and Windows operating systems. For Sun Solaris operating systems, VERITAS offers such software. These software products enable you to upgrade XP firmware while the application service is running.

Network communications links between the dispersed data centers must be redundant and physically routed differently. This prevents the “backhoe issue,” that is, where all links between data centers are cut together. This is especially important, since the cluster is more vulnerable to “split brain” syndromes. A split brain syndrome is where both data centers’ systems form new clusters which could allow access to both copies of the data. This can be prevented with physically separated network links and redundant network components. Cluster Extension XP allows you to configure the failover behavior in such a way that the application service startup procedure will be stopped if none of the remote cluster members can be reached. The default configuration of Cluster Extension XP expects the cluster software to deal with the “split brain” syndrome.

Since the disk array stores your most valuable data, this data must get across to the remote disk array. At least four Continuous Access XP links must be available when the disk arrays are connected directly and are configured for bidirectional takeover. For extended distances, extender components must be purchased. These components are able to bundle Continuous Access XP links. At least two links are necessary to provide redundancy and protection against single points of failure. Although communications links can cover considerable distances, each network segment must be extended to the dispersed data center in order to maintain a heartbeat among all servers.

Recommendation Use four systems to give local application service failover among local cluster systems priority over remote, more time-consuming failover procedures. When failing over, Cluster Extension XP must reconfigure the disk arrays to change the mirroring direction. This takes more time than just checking for the correct disk array disk states. On the remote site, two systems should be available in the case the failover system experiences a hardware or power failure.

The preferred Cluster Extension XP configuration is depicted on [page 29](#).

Caution *Cluster Extension XP works with only one system at each location, with a single I/O path between the server system and the disk array and a single link in each direction between disk arrays.*

However, those configurations are not considered highly available, nor are they disaster tolerant. Therefore, Cluster Extension XP configurations with single points of failure are not supported by HP.

Cluster Extension XP execution

Cluster Extension XP requires cluster software to automatically fail over and fail back among systems on a local site or between sites. Cluster Extension XP must manipulate the application startup process before disk array disks are activated. Cluster Extension XP, therefore, must be integrated as first resource (in the order of resources). To activate Continuous Access XP paired disk devices, the paired disk devices must be in read/write mode. Continuous Access XP disks are usually in read/write mode on the primary disk only; the secondary disk is in read-only mode. In case of a failover, the direction of the mirrored pair is changed by Cluster Extension XP automatically. In case of a disaster, the disk array can have several different states for disks in a RAID Manager XP device group. Cluster Extension XP decides whether those disks can be activated.

Cluster Extension XP must be installed on any server in the cluster that can run the application service in the cluster.

Cluster Extension XP stores information about the application environment in an internal object database and uses RAID Manager XP to gather information about the state of the associated disk pairs. The information about the configured disk array environment and failover behavior is transferred either directly by the cluster software or by gathering from the user configuration file.

The internal object database provides Cluster Extension XP with knowledge about supported parameters, their formats, and default values.


Disk array disk states are stored in an internal object database and a rule engine is used to process those disk states. The rule engine matches current disk states and configuration parameters with a defined rule, stores it in the database, and invokes predefined actions. Those actions prepare the disk array disks to be activated, or it stops the application service startup process if the matching rule requires it to do so.

Continuous Access XP and RAID Manager XP

Continuous Access XP provides remote copy functionality for the disk arrays. Disk arrays can be mirrored to many different remote disk arrays.

Cluster Extension XP does not support two disk arrays as either primary or secondary disk arrays. Cluster Extension XP supports configurations where two (or more) disk arrays use one remote disk array as the failover site. In those cases, the disk array configuration can be considered as a logical one-to-one configuration.

An example of a supported configuration is depicted on [page 43](#).



Supported XP disk array configuration

To control Continuous Access XP-mirrored disks from a server, RAID Manager XP must be installed on the server. A special disk, called a *command device*, must be configured to control the paired disks. The special disk must not be part of Microsoft Cluster Service resources and cannot be paired. The command device, which is identified by a “CM” appended to the emulation type, can be assigned to a 36-Mbyte or greater CVS volume. RAID Manager XP uses the command device to communicate with the disk array controller (DKC).

Using Continuous Access XP Extension, consistency groups can be configured. Consistency groups are units in which the disk array keeps data consistent among paired disks.

Continuous Access XP links are unidirectional links. For disaster tolerant configurations, two links must be provided in each direction. Both sender (RCP) and receiver (LCP) ports must be configured on each redundant I/O board used for Continuous Access XP.

Continuous Access XP offers two modes of replication:

- synchronous replication
- asynchronous replication

Synchronous replication

Using synchronous mode, all write requests from the server are first transferred to the remote disk array. After each I/O has been mirrored in the cache area of the remote array, it is acknowledged to the local disk array. The write request is then acknowledged to the server.

Synchronous replication modes can be configured in the following fence levels:

NEVER	Allows write requests even if the request cannot be replicated to the remote disk array. If a write request cannot be replicated to the remote disk array, the area on the disk is marked in a bitmap table and transferred after a resynchronization request has been ordered.
STATUS	This fence level is not supported by Cluster Extension XP.
DATA	Prohibits write requests immediately if a link failure or disk failure occurs. The local disk array cannot replicate data to the remote disk array. Fence level DATA provides data concurrency at any time.

The preceding fence levels provide data integrity on a per disk basis, so a failure affecting a single disk pair does not lead to a halt of the replication activities of non-affected disk pairs.

Synchronous replication can affect the performance of the system if the distance between the disk arrays is significant.

Asynchronous replication

Continuous Access XP Extension offers a unique feature to replicate data asynchronously.

To keep replicated data consistent among two disk arrays, any incoming write request is ordered and numbered. The write request is then acknowledged to the server, offering the fastest response time for remote mirroring. Each write request is transferred to the remote disk array asynchronously. The remote array orders all write requests before they are destaged to the disk, keeping data consistent.

Asynchronous replication offers excellent performance for remote mirroring and provides data consistency on a group of disks (consistency groups) level.

RAID Manager XP instances

A RAID Manager instance is necessary to control pair operations and to gather disk array status information.

The RAID Manager XP instance numbers used for the **RaidManagerInstances** object must be the same among all systems using Cluster Extension XP.

Several RAID Manager XP instances can be configured to provide additional redundancy. Cluster Extension XP switches to the next available instance when an instance becomes unavailable.

The RAID Manager XP instances should be running at all times to provide the fastest failover capability. Cluster Extension XP provides scripts to include the RAID Manager XP startup procedure in the system startup file

(for example, **/etc/inittab**). However, Cluster Extension XP starts the configured RAID Manager XP instances if it cannot find any running instance.

Quorum service

The Cluster Extension XP quorum service employs static RAID Manager API calls and therefore is not dependent on a RAID Manager instance.

RAID Manager XP device groups

A single device group must be configured for a service group (VCS), a resource group (HACMP), a cluster group (Microsoft Cluster Service), or a package (SG-LX). This device group must include all disks being used for the application service.

The device group is the unit in which the failover/failback operation is being carried out. A device group can contain several volume groups.

User configuration file and Cluster Extension XP objects

Objects define the disk array environment and failover/failback behavior. Objects can be customized in the user configuration file or directly in the cluster software.

The user configuration file

Cluster Extension XP uses the user configuration file to gather application service-specific information. This file describes the dependencies between application services and RAID Manager XP device groups in one file for all application services in the cluster. This file must be copied to all nodes that use Cluster Extension XP.

The user configuration file must be placed in the configuration directory:

Linux **/etc/opt/hpclx/conf**
UNIX

Windows By default, this location is defined as this value:
%ProgramFiles%\Hewlett-Packard\Cluster Extension XP\conf

Related information “Basic configuration example” ([page 89](#))
“Creating and configuring the user configuration file” ([page 197](#))

HACMP

The **UCF.cfg** file is required for IBM HACMP. A single **UCF.cfg** file must be maintained and copied to all systems using Cluster Extension XP. The **UCF.cfg** includes a “common” section to configure the Cluster Extension XP environment and an “application” section to configure the application service-dependent failover/failback behavior. The application section is a multitag component; the **APPLICATION** tag and application-related objects can appear numerous times in the **UCF.cfg**.

Related information “User configuration file for HACMP” ([page 103](#))

Microsoft Cluster Service

Cluster Extension XP integration with Microsoft Cluster Service does not require a user configuration file when the standard environment for Cluster Extension XP is used. The Cluster Extension XP objects that are integrated with Microsoft Cluster Service are configurable as resource private properties in the cluster software.

Related information “Configuring Cluster Extension XP resources” ([page 119](#))

VCS

Cluster Extension XP integration with VERITAS Cluster Server does not require a user configuration file when the standard environment for Cluster Extension XP is used. The Cluster Extension XP objects that are integrated with VERITAS Cluster Server are configurable as resource attributes in the cluster software.

Related information “Configuring the Cluster Extension XP resource” ([page 156](#))

SG-LX

An environment configuration file is required for Serviceguard. The file must reside in the same directory as the package control file and is identified by the package name:

*package_name***_clx.env**

The **APPLICATION** tag is required, although no value is required.

Related information “Configuration of the Cluster Extension XP environment” ([page 176](#))

File structure

The configuration file comprises a common section and application sections. These sections are distinguished by control tags. Cluster Extension XP uses the following objects as control tags:

- **COMMON**
- **APPLICATION**

Objects have one of the following formats:

tag	a definition of an object, for example, COMMON or APPLICATION .
integer	a number, for example, a timeout value.
string	a name, which can include alphabetic and numeric characters and underscores, for example, an application startup value.
list	a list of space-separated strings, for example, a list of host names (lists of numbers are stored as lists of strings).

Text that is a comment starts with the pound (#) character and continues until the end of the line. Comments can start on a new line or be part of a line specifying an object.

Specifying object values

When using the default configuration, you must provide values for these five objects:

- DeviceGroup** ([page 84](#))
- DC_A_Hosts** ([page 84](#))
- DC_B_Hosts** ([page 84](#))
- RaidManagerInstances** ([page 86](#))
- XPSerialNumbers** ([page 88](#))

You do not need to change the default settings unless you want to change the degree of protection for your paired disks. If you change an object, you may need to change additional objects as well. For example, if you change the **FenceLevel** object to **DATA**, you might need to change the **DataLoseMirror** object also.

Objects are supported according to the requirements or capabilities of the cluster software, as listed in table 1 ([page 71](#)).

Table 1. Cluster Extension XP supported objects

Name	Page	CLI	HACMP	MS Cluster service	VCS	SG-LX
COMMON	73	•	•	•	•	•
LogDir	73	•	•	•	•	•
LogLevel	73	•	•	•	•	•
SearchObject	74		•			
VcsBinPath	74				•	
APPLICATION	76	•	•	•	•	•
ApplicationDir	76	•	•	•	•	
ApplicationStartup	77	•	•	•	•	•
AsyncTakeoverTimeout	79	•	•	•	•	•
AutoRecover	80	•	•	•	•	•
BCEnabledA	81	•	•	•	•	•
BCEnabledB	81	•	•	•	•	•
BCMuParamA	81	•	•	•	•	•
BCMuParamB	81	•	•	•	•	•
BCResyncEnabledA	81	•	•	•	•	•
BCResyncEnabledB	82	•	•	•	•	•
BCResyncParamA	82	•	•	•	•	•
BCResyncParamB	82	•	•	•	•	•
DataLoseDataCenter	82	•	•	•	•	•
DataLoseMirror	83	•	•	•	•	•
* DC_A_Hosts	84	•	•	•	•	•
* DC_B_Hosts	84	•	•	•	•	•

Table 1. Cluster Extension XP supported objects (Continued)

Name	Page	CLI	HACMP	MS Cluster service	VCS	SG-LX
<i>(continued)</i>						
* DeviceGroup	84	•	•	•	•	•
FastFailbackEnabled	85				•	
FenceLevel	85	•	•	•	•	•
Filesystems	85	•	•			
PostExecCheck	86	•	•	•	•	•
PostExecScript	86	•	•	•	•	•
PreExecScript	86	•	•	•	•	•
* RaidManagerInstances	86	•	•	•	•	•
ResyncMonitor	87		•	•	•	•
ResyncMonitorAutoRecover	87		•	•	•	•
ResyncMonitorInterval	87		•	•	•	•
ResyncWaitTimeout	88	•	•	•	•	•
Vgs	88	•	•	•	•	•
* XPSerialNumbers	88	•	•	•	•	•

LEGEND

* Required

• Supported

COMMON section objects

The common part is used to set the environment of Cluster Extension XP.

The **COMMON** tag is a single-tag; it can appear in the configuration file only once. The common object does not require any value.

Objects of the type common can only appear once. Those objects must be placed after the **COMMON** tag in the configuration file.

If the default values fit your environment, there is no need to specify them in the file.

COMMON

Format tag

Description Distinguishes between general (common) and application-specific objects.

LogDir

Format string

Description (*Optional*) Defines the path to the Cluster Extension XP log file.

Default value *Linux/UNIX*
/var/opt/hpclx/log

Windows
%ProgramFiles%\Hewlett-Packard\Cluster Extension XP\log

LogLevel

Format string

Description (*Optional*) Defines the logging level used by Cluster Extension XP.

<i>Valid values</i>	error (<i>default</i>)	Logs only error messages for events that are nonrecoverable.
	warning	Logs error messages and warning messages for events that are recoverable.
	info	Logs error messages, warning messages, and additional information, such as disk status.
	debug	Logs error messages, warning messages, info messages, and messages that report on execution status, useful for troubleshooting.

SearchObject *HACMP only*

Format string

Description (*Optional*) Searches for the application service if the user configuration file specifies multiple applications. This object is not used for VCS, Microsoft Cluster Service, or SG-LX.

Default value **Vgs**

VcsBinPath *VCS only*

Format string

Description (*Optional*) Defines the path to the VCS binaries. This object is not used for Microsoft Cluster Service, SG-LX, or HACMP.

Default value **/opt/VRTSvcs/bin**

APPLICATION section objects

The application part defines the failover and failback behavior of Cluster Extension XP for each application service. **APPLICATION** is a multitag that can appear in the configuration file for each application service using Cluster Extension XP.

The **APPLICATION** object requires the name of the application service as its value. The objects specified after an **APPLICATION** tag must appear only once per application. As with the common part objects, the application part objects have predefined default values.

Cluster Extension XP also uses the following rules to define objects:

- If you use the default value, you do not have to specify the object.
- Cluster Extension XP uses objects depending on the setting of other objects. For example, if you set the **FenceLevel** object to **DATA**, Cluster Extension XP uses the values specified for the **DataLoseMirror** or **DataLoseDataCenter** object. However, these objects are ignored if the **FenceLevel** object is set to **NEVER**.
- The pre-execution and post-execution functions in Cluster Extension XP will not be processed if the associated object values are empty. (This is the default setting.)

CLI
HACMP
SG-LX

To set **APPLICATION** object values, use the user configuration file.

VCS

Use the VCS GUI to set **APPLICATION** object values.

Microsoft Cluster Service

To set **APPLICATION** object values, use the Microsoft Cluster Service Cluster Administrator GUI.

APPLICATION

Format tag

Description Distinguishes between general and application-specific objects. Specify the name of the application service. The format of its value is equivalent to a string value.

SG-LX

For Serviceguard, the tag is required; however, specifying a value is not necessary.

ApplicationDir

Format string

Description Specifies the directory where Cluster Extension XP searches for application-specific files, such as the force flag or online file.

If **ApplicationDir** is set to a nonexistent drive and **PairResyncMonitor** is not enabled, Cluster Extension is unable to create the online file and cannot put the resource online.

SG-LX

The value of **ApplicationDir** is derived from the package control file location.

Windows

If **ApplicationDir** is not set, Cluster Extension uses the local **%HPCLX_PATH%** values as defined in the registry.

Default values

Linux

UNIX

/etc/opt/hpplx

Windows

%HPCLX_PATH%

Files *resource_name.createplitbrain*
resource_name.forceflag
resource_name.online

If specified in a user configuration file, *resource_name* is the value of the **APPLICATION** tag; otherwise, *resource_name* is the value of the Cluster Extension XP resource name.

ApplicationStartup

Format string

Description (*Optional*) Specifies where a cluster group should be brought online.

The **ApplicationStartup** object can be customized to determine whether an application service starts locally or is transferred back to the remote data center (if possible) to start directly without waiting for resynchronization. This object is used only if an application service has already been transferred to the secondary site and no recovery procedure has been applied to the disk set (the disk pair has not been recovered and is not in **PAIR** state). This process is considered a failback attempt without prior disk pair recovery.

Cluster Extension XP can detect the most current copy of your data based on the disk state information. If Cluster Extension XP detects that the remote XP disk array has the most current data, it orders a resynchronization of the local disk from the remote disk, or it stops the startup process to enable the cluster software to fail back to the remote XP disk array.

If a resynchronization is ordered, Cluster Extension XP monitors the progress of the copy process. If the application service was running on a secondary XP disk array without replication link, a large number of records may need to be copied. If the copy process takes more time than the configured application startup timeout, the application startup will fail.

Microsoft Cluster Service

If the **ApplicationStartup** resource property is set to **FASTFAILBACK** and the **FailoverThreshold** value is set to a number higher than the current number of clustered systems for the resource group, the resource group will restart on configured nodes until one of the following conditions is met:

- The resource is brought online in the remote data center.
- The resource failed because the **FailoverThreshold** value has been reached.
- The resource failed because the **FailoverPeriod** timeout value has been reached.

Caution

Disable subsequent automated failover procedures for recovery failback operations.

Valid values

FASTFAILBACK (default)

The cluster group will be brought online in the remote data center (if possible) without waiting for resynchronization. The application startup process will be stopped locally and Cluster Extension XP reports a data center error. Depending on the cluster software, the application service cannot start on any system in the local data center and the cluster software will transfer the application service back to the remote data center. Use this value to provide the highest application service uptime. Depending on the value configured for the **AutoRecover** object, Cluster Extension XP will attempt to update the former primary disk based on the secondary disk and swap the personalities of the disk pair so that the local disk will become the primary disk.

In a two-node cluster, this process will not work because the target failback system would not be available. In this case, the application service must be started manually, or the **ApplicationStartup** object should be set to **RESYNCWAIT**.

In a CLX for Microsoft Cluster Service (MSCS) integration, Cluster Extension XP can detect when there is no target failback system available in the remote data center. In this case, Cluster Extension XP will behave as if the **ApplicationStartup** resource property is set to **RESYNCWAIT**.

RESYNCWAIT Online local, cluster group must wait until the disk status is **PAIR**. Cluster Extension XP will initiate a resynchronization of the local disk based on the remote disk. The copy process will be monitored. If no copy progress was made after a monitoring interval expired, the copy process is considered failed and Cluster Extension XP returns a global error. If **RESYNCWAIT** has been specified for the **ApplicationStartup** object, the **ResyncWaitTimeout** object must be specified, in case Cluster Extension XP should wait for resynchronization changes for more or less than 90 seconds, which is the default.

AsyncTakeoverTimeout

Format integer

Description (*Optional*) Specifies the **horctakeover** command timeout in seconds. Must be adjusted based on disk mirroring link speed.

This object is used only if the **FenceLevel** object value is **ASYNC**.

The takeover operation for fence level **ASYNC** (Continuous Access XP Extension) offers the option to stop the data transfer process after a specified time value. This is used to allow access to the remote copy if the data transfer process has been stopped due to a Continuous Access XP-link failure. All data that has been copied up to the moment the timeout value has been reached is consistent and available to access at the secondary site.

Caution

Measure or calculate the full XP disk array cache copy time to use the gathered information for the **AsyncTakeoverTimeout** object. After a takeover command has been invoked, Continuous Access XP Extension copies the side file area residing in the XP disk array cache to the site where the takeover command has been issued (the secondary disks). The side file area cannot exceed the installed cache size. The maximum time for the **AsyncTakeoverTimeout** object is the time to fully copy the amount of cache size data. The takeover timeout value is used to terminate the copy process to provide access to the secondary disks, for example, if all links or the primary XP disk array are unavailable to copy the side file area. The copy time depends on the performance of the Continuous Access XP link between your sites. The takeover or resynchronization operation could take longer than the timeout value for application service startup in the cluster software. The application service startup might fail in this case. However, the takeover or resynchronization command will continue in the background.

Default value **1800** (default)

AutoRecover

Format string

Description (Optional) Recovers a suspended or deleted disk pair when the resource is brought online at application service startup time.

If the **AutoRecover** object is set to **YES**, Cluster Extension XP will try to resynchronize the remote disk at application startup time. Cluster Extension XP will ignore the return code of the resynchronization command and allow access to the disk ensuring highest application availability.

If the resynchronization attempt fails, Cluster Extension XP will not fail. The internal logic will first apply the concurrency and consistency rules to allow access to the disk set.

If you configure fence level **DATA** for the device group and set the **FenceLevel** object to **DATA**, the **AutoRecover** object will change Cluster Extension XP's behavior. Cluster Extension XP will attempt to reestablish the **PAIR** state and wait for the **PAIR** state before it allows access to the disk. If the resynchronization or takeover process fails, Cluster Extension XP returns a global error.

Valid values **YES** (*default*)
NO

BCEnabledA

Format string
Description (*Optional*) Enables rolling disaster protection for data center A.
Valid values **YES**
NO (*default*)

BCEnabledB

Format string
Description (*Optional*) Enables rolling disaster protection for data center B.
Valid values **YES**
NO (*default*)

BCMuListA

Format list
Description (*Optional*) Space-separated list defines the MU number of the Business Copy XP disk pairs in data center A.

BCMuListB

Format list
Description (*Optional*) Space-separated list defines the MU number of the Business Copy XP disk pairs in data center B.

BCResyncEnabledA

Format string
Description (*Optional*) Enables automatic resynchronization of Business Copy XP disk pairs in data center A. The automatic resynchronization function is supported only when the split BC pair is located in the same data center where Cluster Extension XP is started.
Valid values **YES**
NO (*default*)

BCResyncEnabledB

<i>Format</i>	string
<i>Description</i>	(Optional) Enables automatic resynchronization of Business Copy XP disk pairs in data center B. The automatic resynchronization function is supported only when the split BC pair is located in the same data center where Cluster Extension XP is started.
<i>Valid values</i>	YES NO (default)

BCResyncMuListA

<i>Format</i>	list
<i>Description</i>	(Optional) Space-separated list defines the MU number of the Business Copy XP disk pairs in data center A.

BCResyncMuListB

<i>Format</i>	list
<i>Description</i>	(Optional) Space-separated list defines the MU number of the Business Copy XP disk pairs in data center B.

DataLoseDataCenter

<i>Format</i>	string
<i>Description</i>	(Optional) Specifies whether a resource should be brought online while the disk pair is (or will be) suspended or deleted and there is no connection (CA XP and IP network) to the remote data center.

Used only if the **FenceLevel** object value is **DATA**.

RAID Manager XP is able to access its remote peer to invoke takeover actions for Continuous Access XP device groups. It is also able to invoke a swap-takeover operation of the device group from the secondary site. If no configured remote RAID Manager XP instance replies to a request of the local RAID Manager XP instance (remote status EX_ENORMT), all network connections between the local and the remote data center are considered *DOWN*. If the swap-takeover operation leads into a suspended state for the device group, the Continuous Access XP links are considered *DOWN*.

Because redundant networks and Continuous Access XP links are necessary to build a disaster tolerant environment, this situation can be considered as a data center failure. The **DataLoseDataCenter** object is used to allow/prohibit automatic application service startup in this particular case.

The combination of setting the **DataLoseMirror** object to **YES** and the **DataLoseDataCenter** object to **NO** are contradictory.

Valid values **YES** (*default*)
NO

DataLoseMirror

Format string

Description (*Optional*) Specifies whether a resource should be brought online while the disk pair is suspended or deleted.

Used only if the **FenceLevel** object value is **DATA** and local and remote XP disk status information can be gathered. If the remote XP disk state information is not available (remote state **EX_ENORMT**), the setting of the **DataLoseDataCenter** object will be used.

Depending on the value configured for the **AutoRecover** object, Cluster Extension XP will attempt to recover the **PAIR** state for the device group. Cluster Extension XP waits until the **PAIR** state has been established. If this operation fails, Cluster Extension XP returns a global error. Because the **DATA** fence level ensures no loss of concurrency, manual intervention is required to recover the **PAIR** state. The **PAIR** state must be reestablished for all disks in the device group before you can start the application service.

The combination of setting the **DataLoseMirror** object to **YES** and the **DataLoseDataCenter** object to **NO** are contradictory.

Valid values **YES**
NO (*default*)

DC_A_Hosts *Required*

Format list

Description Space-separated list defines the cluster nodes in data center A.

VCS

This object is a string-vector element. Add a new element to the list for each system name.

DC_B_Hosts *Required*

Format list

Description Space-separated list defines the cluster nodes in data center B.

VCS

This object is a string-vector element. Add a new element to the list for each system name.

DeviceGroup *Required*

Format string

Description RAID Manager XP device group, containing the application service disk set.

Files *Linux*

UNIX
/etc/horcmX.conf

Windows

drive:\winnt\horcmX.conf
%system_root%\horcmX.conf

where X is the RAID Manager XP instance number.

FastFailbackEnabled *VCS only*

Format string

Description (*Optional*) Disables VCS service groups for the data center. This allows transferring the service group back to the remote data center immediately. To allow this operation, the VCS configuration file (**main.cf**) will be write enabled and saved later.

The service group will be disabled for all systems contained in either the **DC_A_Hosts** object or **DC_B_Hosts** object. Then, the VCS configuration file will be saved (dumped).

Valid values **YES** (*default*)
NO

FenceLevel

Format string

Description (*Optional*) The **FenceLevel** object specifies the fence level configured for the device group. Cluster Extension XP checks whether the current fence level reported by the XP disk array is the same as the configured (expected) fence level. This object is also used to make sure your configurations are supported based on consistency considerations. Different failover and recovery procedures are used for different fence levels.

If you change the **FenceLevel** object value, also review the values of these objects:

DataLoseMirror ([page 83](#))
DataLoseDataCenter ([page 82](#))
AsyncTakeoverTimeout ([page 79](#))

Valid values **DATA**
NEVER (*default*)
ASYNC

Filesystems *CLI and HACMP only*

Format list

Description Space-separated list of file systems.

PostExecCheck

<i>Format</i>	string
<i>Description</i>	(Optional) The PostExecCheck object is used to configure Cluster Extension XP to gather XP disk pair status information after the takeover procedure. That information will be passed to the post-executable. In case of a remote data center failure, it could be time consuming to gather that information, especially if your post-executable does not need any XP status information. The arguments passed to the post-executable will include only the local disk status if the PostExecCheck object is set to NO . See “RAID Manager XP configuration” (page 92).
<i>Valid values</i>	YES NO (default)

PostExecScript

<i>Format</i>	string
<i>Description</i>	(Optional) Specifies an executable with its full path name to be invoked after the takeover action or failover procedure.

PreExecScript

<i>Format</i>	string
<i>Description</i>	(Optional) Specifies an executable with its full path name to be invoked before the takeover action or failover procedure.

RaidManagerInstances

Required

<i>Format</i>	list
<i>Description</i>	A space-separated list of RAID Manager XP instances Cluster Extension XP can use to communicate with the disk array. The instance numbers must be the same among all cluster systems. Cluster Extension XP can alternate between the specified instances.

VCS

This object is a string-vector element. Add a new element to the list for each system name.

Files *Linux*
 UNIX
 /etc/horc*mX*.conf

Windows
%systemroot%\horcm*X*.conf

where *X* is the RAID Manager XP instance number.

ResyncMonitor

Format string

Description (*Optional*) Starts the pair/resync monitor to monitor the disk pair status and resynchronize disk pairs if the **ResyncMonitorAutoRecover** attribute is set to **YES**.

Valid values **YES** (*default: Microsoft Cluster Service*)
 NO (*default: HACMP; SG-LX; VCS*)

ResyncMonitorAutoRecover

Format string

Description (*Optional*) Automatically recovers disk pairs states if the disk pairs are monitored by the pair/resync monitor.

Valid values **YES**
 NO (*default*)

ResyncMonitorInterval

Format integer

Description (*Optional*) Specifies the monitor interval (in seconds) that the pair/resync monitor checks the disk pair status.

Default value **60**

ResyncWaitTimeout

Format integer

Description (*Optional*) Specifies the timeout value (in seconds) for a disk pair resynchronization. It may take some time to resynchronize disks. The timer times out if there is no change in the percentage value of the copy status for the device group in the specified time interval. The timeout value is used if the **ApplicationStartup** object is set to **RESYNCWAIT**.

Default value **90**

Vgs

CLI and HACMP only

Format list

Description List of volume groups

XPSerialNumbers

Required

Format list

Description A space-separated list of at least two serial numbers must be specified: the serial numbers of the primary and secondary XP disk arrays. Cluster Extension XP checks whether the local disk array is contained in this list. Serial numbers of the disk arrays of the connected cluster nodes (at least two).

VCS

This object is a string-vector element. Add a new element to the list for each system name.

Basic configuration example

The following is an example of a basic **UCF.cfg** file.

```
#/etc/opt/hpclx/conf/UCF.cfg
#This is the Cluster Extension XP User Configuration File (UCF.cfg).
#The COMMON tag specifies the configuration for the
#Cluster Extension XP core environment
COMMON
LogLevel          info                #default (not necessary)
APPLICATION       sap                 #the application service
Vgs               sapdatavg saptmpvg  #the volume groups (not necessary)
Filesystems       /sapdata /saptmp    #the filesystems
DeviceGroup       sapdg               #RM dev group for the app service
RaidManagerInstances 22              #RM instance number for dev group
DC_A_Hosts        host1a host2a      #Data center A
DC_B_Hosts        host3b host4b      #Data center B
```

