



Configuring HP StorageWorks XP Cluster Extension for SteelEye LifeKeeper

January 2011

HP is a registered trademark of Hewlett-Packard Company.

©2011 SIOS Technology Corp.
San Mateo, CA U.S.A.
All Rights Reserved

Table of Contents

- Introduction 1
- Document Contents 1
- LifeKeeper Documentation 1
- HP Documentation 1
- Requirements 2
 - Hardware Requirements 2
 - Software Requirements 2
- Configuring LifeKeeper 3
 - Specific Configuration Considerations for LifeKeeper Recovery Kit for XP Cluster Extension..... 3
 - Sample UCF.cfg..... 5
 - Recommended UCF.cfg Settings for LifeKeeper 5
 - Configuring a LifeKeeper Hierarchy..... 7
- Troubleshooting Hints and Tips 9
- Appendix A 10
 - Example of an XP RAID Manager Instance Configuration File 10
- Appendix B..... 11
 - Mounting Multipath Device and Mapping DeviceGroup LUNs to the Multipath Device 11

Introduction

LifeKeeper for Linux provides clustering support for applications residing on top of file systems that are being replicated using HP StorageWorks XP Continuous Access software by integrating with HP StorageWorks XP Cluster Extension.

Document Contents

This document explains the following topics:

- [LifeKeeper Documentation](#). A list of all LifeKeeper for Linux documentation and where the information is available.
- [HP Documentation](#). A list of all HP documentation and where the information is available.
- [Requirements](#). Before you can install and set up this software, your server must meet certain hardware and software requirements.
- [Configuring HP StorageWorks XP Cluster Extension for LifeKeeper](#). To ensure that your LifeKeeper product provides the protection and flexibility you require, you need to plan your configuration. To appropriately plan your configuration, you must understand your network configuration, interface selection, user system setup, hierarchy options and the HP XP configuration tasks. In addition to planning your configuration, this section also includes configuration examples and the specific tasks required to configure HP StorageWorks XP Cluster Extension for LifeKeeper.
- [Troubleshooting Hints and Tips](#). This section provides a list of hints, tips and messages that you may encounter during the process of create, extend, failover, local recovery, switchover, and disaster recovery.

LifeKeeper Documentation

The following LifeKeeper product documentation is available from SIOS Technology Corp.:

- *LifeKeeper for Linux Release Notes*
- *LifeKeeper for Linux Online Product Manual* (available from the Help menu within the LifeKeeper GUI)
- *LifeKeeper for Linux Planning and Installation Guide*

This documentation, along with documentation associated with optional LifeKeeper recovery kits, is available on the SIOS Technology Corp. website at:

<http://us.sios.com/support>

HP Documentation

For detailed information and instructions on HP StorageWorks XP Raid Manager, HP StorageWorks XP Continuous Access and HP StorageWorks XP Cluster Extension software, refer to the following product manuals:

- [HP StorageWorks XP Raid Manager User's Guide](#)
- [HP StorageWorks XP24000/XP20000 Continuous Access Software User's Guide](#)
- [HP StorageWorks XP Cluster Extension Software Administrator Guide](#)

Requirements

For proper operation and support, ensure that your configuration meets the following requirements:

Hardware Requirements

- **Servers.** This solution requires a cluster configured with HP StorageWorks XP Continuous Access Software and HP StorageWorks XP Cluster Extension for Linux Software. It must also be configured in accordance with the LifeKeeper requirements described in the *LifeKeeper Online Product Manual* and the *LifeKeeper for Linux Release Notes*, which are shipped with the product media.
- **System Architectures.** x86 (Intel and AMD) 32-bit and x86 (Intel and AMD) 64-bit.
- **Data Storage.** This package requires two HP StorageWorks XP24000 Disk Arrays configured by HP StorageWorks XP Continuous Access Software and HP StorageWorks XP Cluster Extension Software.

Software Requirements

- **Operating System.** SLES 10 SP3, SLES 11 or Red Hat Enterprise Linux 5.4 or 5.5.
- **LifeKeeper Software.** The same version of LifeKeeper core software and any patches plus the LifeKeeper Recovery Kit for XP Cluster Extension, the LifeKeeper DMMP Recovery Kit and the LifeKeeper Quorum/Witness Server Package (optional) must be installed on each server. Please refer to the “*LifeKeeper for Linux Release Notes*” for specific LifeKeeper requirements.
- **Other Software.** Requires an activated installation of DMMP, HP StorageWorks XP RAID Manager and HP StorageWorks XP Cluster Extension on the servers and HP StorageWorks XP Continuous Access Software on the disk arrays.

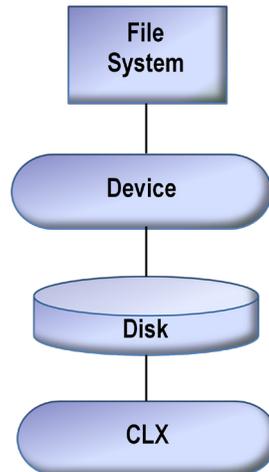
Consult the “*LifeKeeper for Linux Release Notes*” or your sales representative for the latest release compatibility and ordering information.

Consult the HP Single Point of Connectivity Knowledge (SPOCK) website for detailed information about supported HP StorageWorks product configurations at <http://h20272.www2.hp.com>.

Configuring LifeKeeper

The LifeKeeper Recovery Kit for XP Cluster Extension supports HP StorageWorks XP Continuous Access Software by providing protection for applications running on servers attached to XP storage arrays.

The file system hierarchy created with the LifeKeeper Recovery Kit for XP Cluster Extension is:



The CLX resource will be the terminal node in the hierarchy. This allows the CLX resource to manipulate the state of the replication on the XP array and to recognize that a disk (LUN) is shared when the disk IDs are different.

Specific Configuration Considerations for LifeKeeper Recovery Kit for XP Cluster Extension

This section contains information that should be considered **before** LifeKeeper is configured for a particular application, device group or LUN.

Before configuring the LifeKeeper Recovery Kit for XP Cluster Extension, the following steps should be taken:

1. Install operating system and the following software packages:

HP Software

- HP StorageWorks XP RAID Manager
- HP StorageWorks XP Cluster Extension for Linux Software
- Device Mapper Multipath Enablement Kit for HP

SIOS Technology Corp. Software

- LifeKeeper Core 7.2.0 or later and Installation Support 7.2.0 or later
- LifeKeeper DMMP Recovery Kit 7.2.0 or later
- LifeKeeper Quorum/Witness Package 7.2.0 or later (Optional)
- LifeKeeper Recovery Kit for XP Cluster Extension 7.2.0 or later

2. Confirm Connections

Confirm that the physical connections from the servers to the XP arrays are operational and that DMMP lists the correct paths to the storage.

3. Configure HP StorageWorks XP RAID Manager

- Edit the `services` and `hosts` files on each server that runs instances (`/etc/services` and `/etc/hosts`).

Note: UDP Port Number in the `services` file must match the UDP Port Number in `horcm#.conf`. In the `services` file, add an entry for each instance and add one `clxmonitor` entry for the daemon.

For detailed instructions on editing the `services` and `hosts` files, refer to the [HP StorageWorks XP Raid Manager User's Guide](#).

- Set up and create the XP RAID Manager Instance Configuration File (`/HORCM/etc/horcm.conf`). The configuration file contains the following sections that will need to be edited: `HORCM_MON` section, `HORCM_CMD` section, `HORCM_DEV` section or `HORCM_LDEV` section, `HORCM_INST` or `HORCM_INSTP` section.

Note: The `mkconf` command can be used to create the configuration file. See [Appendix A](#) for an example of a RAID Manager Instance configuration file.

For detailed instructions on setting up the Instance Configuration File, refer to the [HP StorageWorks XP Raid Manager User's Guide](#).

- After setting up the instance configuration files, the instances can be started up. To manually start an instance, run this shell command on each host that runs an instance:

```
/usr/bin/horcmstart.sh [instance_number] [instance_number]...
```

The instances can also be started during bootup. For this to happen, the file `/etc/init.d/raidmgr` needs each instance defined by a “horcm” config file in the “RAIDMGR INSTANCES” list.

Note: If you do not specify an instance number, the command uses the value stored in the `HORCM_INST` environment variable. The default value is 0.

Make sure that when changes are made in the `horcm` file that similar changes are made on every node in the cluster. The RAID Manager instance will then need to be restarted using `horcmshutdown.sh` and `horcmstart.sh` on all nodes in the cluster.

4. Configure HP StorageWorks XP Continuous Access

When properly configured, HP StorageWorks XP Continuous Access with LifeKeeper Recovery Kit for XP Cluster Extension not only provides data protection through replication, but also failover protection for applications. Remote replication is the continuous copying of data from selected virtual disks on a source (local) array to related virtual disks on a destination (remote) array. Applications continue to run while data is replicated in the background. Remote replication requires a fabric connection between the source and destination arrays and a software connection between source virtual disks and destination virtual disks.

Make sure the following tasks have been completed:

- Disk array fibre channel ports have been configured
- Host and host bus adapters that connect the host to the disk array have been installed and

- configured
- Cabling between the arrays and the fibre channel switches meets remote replication requirements
- Host port data replication settings have been selected
- Replication licenses have been installed on each server for each array
- Fabrics have been created and configured
- Host security settings have been configured
- Source virtual disks have been created and presented
- Mirrored pairs have been successfully created

For detailed information and instructions on configuring your disk arrays, refer to the [HP StorageWorks XP24000/XP20000 Continuous Access Software User's Guide](#).

Configure HP StorageWorks XP Cluster Extension

The Cluster Extension configuration file (`UCF.cfg`) defines clustering rules and gives information on the arrays, instances and device groups, providing the status of the storage and the server cluster. To configure HP StorageWorks XP Cluster Extension, the following must be done:

- Edit the user configuration file (`/etc/opt/hpplx/conf/UCF.cfg`).
Note: Instance # in `UCF.cfg` must match `horcm` instance file and `DeviceGroup` in `UCF.cfg` is defined in `horcm` instance file (`/etc/horcm#.conf`). All entries must be made on every node in the cluster.
- Create and configure an XP Cluster Extension resource.
- Configure the pair/resync monitor.
 - Activate the pair/resync monitor.

Sample UCF.cfg

COMMON

LogLevel error

APPLICATION testA

XPSerialNumbers 78493 78502

RaidManagerInstances 101

DeviceGroup testA

DC_A_Hosts dogwood magnolia

DC_B_Hosts minute-maid

FenceLevel never

ResyncMonitor yes

ResyncMonitorAutoRecover yes

ResyncMonitorInterval 10

AutoRecover yes

ApplicationStartup resyncwait

Recommended UCF.cfg Settings for LifeKeeper

APPLICATION

Application is the same as the DeviceGroup - it can have the same name or a different name

XPSerialNumbers

Serial numbers of the arrays

RaidManagerInstances

Instance number from the horcm file which provides the link between the two files

DeviceGroup

This variable is the same as DeviceGroup in horcm

FenceLevel

This variable determines the behavior of a mirror in failure scenarios -- must match the level configured on the array -- `pairdisplay` will give this information

Valid FenceLevel values:

- never** - "Never" restrict access. In almost all failure scenarios, the data is available and CLX will "handle" any merges using its bitmap. This value places highest priority on ACCESS of the mirror. **Note:** SIOS Technology Corp. recommends that the FenceLevel value be set to **never**.
- data** - Restrict access. This value places highest priority on the consistency of data. This can mean very limited access in failure scenarios.

When FenceLevel value is set to **data**, there are two additional variables which determine whether to allow data access on failure:

DataLoseMirror

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. **[yes/no]**

DataLoseDataCenter

Specifies whether a resource should be brought online while the disk pair is (or will be) suspended or deleted and there is no connection (Continuous Access and IP network) to the remote data center. Please see the HP StorageWorks XP Cluster Extension Software Administrator Guide for more details. **[yes/no]**

When DataLoseMirror is **no**, then DataLoseDataCenter must be **no** as well. When set to **yes/yes**, then data fence level works almost the same as **never** except that when the standby site fails, the primary goes read-only.

ResyncMonitor [yes/no]

This variable should be set to **yes**. In order for LifeKeeper to monitor the application, ResyncMonitor is required. If, for some reason, disabling ResyncMonitor is necessary, this should be done outside of LifeKeeper. **LifeKeeper should be stopped** and then the resync monitor can be stopped.

ResyncMonitorAutoRecover [yes/no]

This variable determines if automatic resync is done after a failure. If set to **yes**, once repair occurs, automatic resync will occur, and then replication will resume. If set to **no**, the mirror will remain in the paused/broken state until manually fixed.

Note: While SIOS Technology Corp. recommends that this parameter be set to **yes**, it could increase the possibility of a "split brain" situation unless either the **LifeKeeper Witness Server** is configured or **Confirm Failover** is enabled. LifeKeeper Witness Server is not an absolute requirement when using the XP Cluster Extension for LifeKeeper. In some situations, it may be desirable to use a different LifeKeeper feature, **Confirm Failover**, when setting up clusters across wide area networks with XP Continuous Access. Configured in the LifeKeeper GUI, this feature blocks failover of resources until an operator confirms that the

failover should be allowed. For complete information on this **Confirm Failover** feature, refer to the *LifeKeeper for Linux Release Notes*.

ResyncMonitorInterval [#]

Specifies the monitor interval (in seconds) that the pair/resync monitor checks the disk pair status. No need to connect this parameter to the HORCM_MON section poll value. Default is 60 seconds.

AutoRecover [yes/no]

This variable determines if automatic resync is done when clxrun is called. If set to **yes**, clxrun will automatically resync the mirror and resume replication. If set to **no**, the mirror will remain in the paused/broken state until manually fixed.

ApplicationStartup [fastfailback/resyncwait]

The ApplicationStartup setting determines what should happen when a mirror is started where the data is not current and CLX can detect the most current copy of the data. If it is determined that the remote system has the most current copy of the data, then depending on how it is configured, it can either wait for the resynchronization of the local disk to complete before it allows access to the mirror or the mirror will fail to start where the mirror will require a resync (so that it can be started where the data is current and access can occur immediately). When **resyncwait** is configured, the clxrun command will wait for the resynchronization to complete. When **fastfailback** is configured and CLX is able to determine that a remote system has the most current copy of the data, then a failback will be forced to the server where a resync is not required and the mirror can start immediately.

Note: Configurations with LifeKeeper will **always** behave in a “**resyncwait**” manner no matter the ApplicationStartup setting.

Note: Make sure that any changes made to the `UCF.cfg` file are also made on every node in the cluster. If changes are made to this file, the resource will need to be taken out of service and then brought back in service in order for the file to be re-read.

The above parameters, along with detailed information and instructions on the configuration of HP StorageWorks XP Cluster Extension, is available within the [HP StorageWorks XP Cluster Extension Software Administrator Guide](#).

Configuring a LifeKeeper Hierarchy

If you have installed your LifeKeeper environment as described in the *LifeKeeper for Linux Planning and Installation Guide*, you should be ready to start and configure the LifeKeeper software on each server in your cluster including the quorum/witness server.

1. Start LifeKeeper.
2. Create communication paths between all servers. Before LifeKeeper protection can be activated, the communication path (heartbeat) definitions must be created within LifeKeeper.

For instructions, see the [Online Help](#) topic “*Creating a Communications Path*” in “*GUI Tasks/Administrator Tasks/Server*.”

3. Mount file system(s) on replicated mirrors. Because monitoring occurs on active/in-service resources (i.e. file systems) only, be sure to mount the file system (and create the resource hierarchy) on the node that is in the P-VOL state and not on a system that is in the S-VOL state (*pairdisplay* displays the states of volumes).

Configuring LifeKeeper

4. Create hierarchy (via GUI).

For instructions on creating hierarchies, see the [Online Help](#) topic “*Creating Resource Hierarchies*” in “*GUI Tasks/Administrator Tasks/Resource/Creating Resource Hierarchies*.”

Note: Storage must be configured and packages installed on all nodes prior to create. It is also required that the RAID Manager instances be running on all nodes prior to create.

5. Extend resource hierarchy to the nodes that are also configured to see storage. These nodes can be in the same site (shared source) or on another site (target). Extending first to the same site and then to the other sites is recommended to keep priorities naturally grouped together.

For instructions on extending resource hierarchies, see the [Online Help](#) topic section “*Extending Resource Hierarchies*” in “*GUI Tasks/Administrator Tasks/Resource*.”

Note: The quorum/witness server should reside at a site apart from the other servers in the cluster to provide the greatest degree of protection in the event of a site failure. For further instructions on the quorum/witness server, please see “Appendix A” in the “[LifeKeeper for Linux Planning and Installation Guide](#)”.

Troubleshooting Hints and Tips

This section provides a list of hints, tips and messages that you may encounter during the process of create, extend, failover, local recovery, switchover, and disaster recovery.

FenceLevel = data (resource OSF on all servers)

In the event of a disaster involving resources with FenceLevel=data, DataLoseMirror=no and DataLoseDataCenter=no, the resource will not be able to come in service on any host in the cluster. These settings indicate to LifeKeeper and CLX that the data integrity is the highest priority for the device group. In the event of a disaster, resources utilizing these settings will fail on all servers to prevent any critical data from being lost or updated from the wrong host.

In a disaster scenario involving resources with these settings, manual intervention is required. The user must manually correct the problem that generated the CA Link failure and/or the Remote DataCenter failure. Once the problem has been corrected and the mirror state has been returned to the P-VOL PAIR state, the LifeKeeper resource should then be brought in service on the correct node.

Create Failure

Verify `UCF.cfg` and `horcm#.conf` files on each server. Verify the RAID Manager instance has been started on every node in the cluster.

Create Occurs with Errors / Extend Failure

Errors in the `UCF.cfg` file can cause the create of a file system resource to contain errors referring to a `pairdisplay` problem and/or can cause the extend to fail with an "ERROR 001008". Verify that the contents and format of the `UCF.cfg` file are correct.

Failover Delay

Depending on the `poll_value` setting in the `HORCM_MON` section of the `horcm` configuration file, there may be a longer than usual delay before resources become active on the failover node.

Partition Device if Needed

The device can be partitioned or the whole device can be used (multipath will have to be forced to reload).

In addition, please refer to the "Troubleshooting" section in the "[LifeKeeper for Linux Online Help](#)".

Appendix A

Example of an XP RAID Manager Instance Configuration File

HORCM_MON

Set IP address to NONE so that horcm will start during bootup
Port must match entry in `/etc/services`

HORCM_CMD

To find command device: `multipath -ll |grep CM`

HORCM_LDEV

DeviceGroup and LUN names are defined here, user preference

HORCM_INST

IP address and port number for target

Example `/etc/horcm101.conf`:

```

HORCM_MON
#IP address    service    poll(10ms)  timeout(10ms)
NONE          52324     1000        500

HORCM_CMD
/dev/mapper/mpathc

HORCM_LDEV
#dev_group    dev name   Serial#    CU:LDEV(LDEV#)
testA        testA_1    78493     00:01

HORCM_INST
# dev_group    ip address  service
testA        172.17.104.95  52324

```

Appendix B

Mounting Multipath Device and Mapping DeviceGroup LUNs to the Multipath Device

To get a list of LUNs associated with DeviceGroup (testA) on the local system, run `pairdisplay`.

```
pairdisplay -l101 -g testA -fwe -l
```

```
Group  PairVol(L/R) (WWN          , LU),Seq#,LDEV#.P/S,Status,Fence,Seq#,P-LDEV# N CTO JID AP EM      E-Seq# E-LDEV# LUN-WWN
testA  testA_1(L) (50060e8015329d00, 0)78493  1.P-VOL PAIR NEVER ,78502  1 - - - 2 -      -      - 60060e8015329d000001329d00000001
```

- DeviceGroup "testA" has one LUN associated with it called "testA_1". The fence level for this DeviceGroup is "NEVER".

To find the multipath device node, find the LUN WWN in multipath.

```
multipath -l | grep 60060e8015329d000001329d00000001
```

```
mpathw (360060e8015329d000001329d00000001) dm-20 HP,OPEN-V
```

- `/dev/mapper/mpathw` is the device to use

To mount partition one (`mpathw_part1`) on the device, issue the mount command.

- `mount /dev/mapper/mpathw_part1 /testA`