



LifeKeeper[®] for Linux v4.7

Data Replication Administration Guide

Please note that this document does not apply to the LifeKeeper Data Replication v5.0.0 or later product. There is a different Administration Guide for LKDR v5.0.0 or later.

May 2006

SteelEye, SteelEye Technology, and LifeKeeper are registered trademarks of SteelEye Technology, Inc.

Adobe Acrobat is a registered trademark of Adobe Systems Incorporation. Apache is a trademark of The Apache Software Foundation. HP and Compaq are registered trademarks of Hewlett-Packard Company. IBM, POWER, DB2, Informix, ServeRAID, Rational and ClearCase are registered trademarks or trademarks of International Business Machines Corporation. Intel, Itanium, Pentium and Xeon are registered trademarks of Intel Corporation. Java is a registered trademark of Sun Microsystems, Inc. Linux is a registered trademark of Linus Torvalds. Microsoft Internet Explorer and Windows are registered trademarks of Microsoft Corporation. MySQL and MaxDB are registered trademarks or trademarks of MySQL AB. Netscape and Netscape Navigator are registered trademarks of Netscape Communications Corporation. NFS is a registered trademark of Sun Microsystems, Inc. Opteron is a trademark of Advanced Micro Devices, Inc. Oracle is a registered trademark of Oracle Corporation and/or its affiliates. PostgreSQL is a trademark of PostgreSQL Global Development Group. Red Flag is a registered trademark of Red Flag Software Co.,Ltd. Red Hat is a registered trademark of Red Hat Software, Inc. SAP is a registered trademark of SAP AG. Sendmail is a registered trademark of Sendmail, Inc. Sun and Solaris are registered trademarks of Sun Microsystems, Inc. SUSE is a registered trademark of SUSE LINUX AG, a Novell business. Sybase is a registered trademark of Sybase, Inc. Other brand and product names used herein are for identification purposes only and may be trademarks of their respective companies.

It is the policy of SteelEye Technology, Inc. to improve products as new technology, components, software, and firmware become available. SteelEye Technology, Inc., therefore, reserves the right to change specifications without prior notice.

To maintain the quality of our publications, we need your comments on the accuracy, clarity, organization, and value of this book.

Address correspondence to:
ip@steeleye.com

Copyright © 2001-2006
By SteelEye Technology, Inc.
Palo Alto, CA U.S.A.
All Rights Reserved

Table of Contents

Document Contents	2
LifeKeeper Documentation	2
Overview	3
LifeKeeper for Linux Data Replication	3
How LifeKeeper for Linux Data Replication Works	3
Resynchronization	4
Features	4
Failover using LifeKeeper for Linux Data Replication	5
LifeKeeper for Linux Data Replication Requirements	7
Hardware Requirements	7
Software Requirements	7
Configuring LifeKeeper for Linux Data Replication	9
LifeKeeper for Linux Data Replication Types	9
Specific Configuration Considerations	10
Configuration Tasks	12
Overview	12
Creating a Data Replication Resource Hierarchy	12
Creating a Data Replication Resource	14
Creating a New Replicated File System	16
Replicating an Existing File System	18
Deleting a Resource Hierarchy	19
Extending Your Hierarchy	20
Extending a Data Replication Resource	23
Extending a New or Existing Replicated File System Resource	24
Unextending Your Hierarchy	26
Testing Your Resource Hierarchy	27
Performing a Manual Switchover from the LifeKeeper GUI	27
LifeKeeper GUI Display During Resynchronization	27
Monitoring the Status of Your NetRAID device	27
Troubleshooting	28
NetRAID Utilities	30
Error Messages	32
Common Error Messages	32
Hierarchy Creation	32
Hierarchy Extension	33
Restore	33
Resource Monitoring	33
Data Replication Error Messages	34
LifeKeeper GUI Related Errors	37

Data Replication Administration Guide

LifeKeeper for Linux Data Replication provides an integrated synchronous data mirroring capability for LifeKeeper for Linux environments. This feature enables LifeKeeper resources to operate in a non-shared storage environment.

Document Contents

This guide contains the following topics:

- [LifeKeeper Documentation](#). A list of LifeKeeper for Linux documentation and where to find them
- [Overview](#). A description of LifeKeeper for Linux Data Replication features and functionality
- [Requirements](#). A description of the hardware and software necessary to properly setup, install, and operate LifeKeeper for Linux Data Replication. Refer to the *LifeKeeper for Linux Planning and Installation Guide* for specific instructions on how to install or remove LifeKeeper for Linux software
- [Configuring LifeKeeper for Linux Data Replication](#). A description of the procedures required to properly configure LifeKeeper for Linux Data Replication
- [Troubleshooting](#). A list of LifeKeeper for Linux error messages including a description for each.

LifeKeeper Documentation

The following LifeKeeper product documentation is available from SteelEye Technology, Inc.:

- *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 SteelEye Technology, Inc. website at:

www.steeleye.com/support/documentation

Overview

LifeKeeper for Linux Data Replication

SteelEye's LifeKeeper for Linux is a high availability solution that ensures the reliability of Linux computing environments using shared storage systems. LifeKeeper for Linux Data Replication offers an alternative to customers who want to use LifeKeeper without the shared storage.

LifeKeeper for Linux Data Replication provides a synchronous data mirroring capability for LifeKeeper protected resources. LifeKeeper protected data is replicated from the currently active server (referenced in this document as the primary server) to a backup server (also known as the target server). LifeKeeper's data replication feature allows data to be reliably, efficiently, and consistently mirrored to remote locations over any IP network, Local Area Network (LAN), or Wide Area Network (WAN).

LifeKeeper for Linux Data Replication provides real time mirroring by employing a synchronous mirroring technique. Synchronous mirroring forwards the primary server writes to the backup site as part of the I/O command processing. Typically, synchronous mirroring awaits remote confirmation before signaling I/O completion to the requesting server. Synchronous mirroring offers the highest level of data protection, ensuring that all copies of the data are always identical. In this mode, data is written simultaneously on the primary and backup servers.

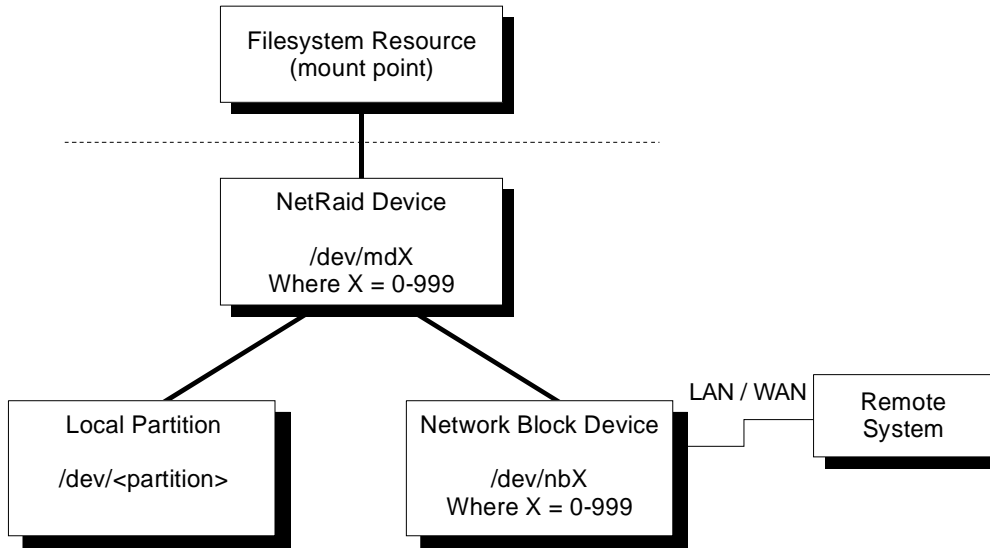
Note: LifeKeeper for Linux Data Replication allows LifeKeeper-supported applications to use NetRAID devices in place of shared storage devices unless otherwise noted in the *Release Notes*.

How LifeKeeper for Linux Data Replication Works

LifeKeeper for Linux Data Replication creates and protects NetRAID devices. A NetRAID device is a data replication device represented as a Linux RAID1 device, which consists of a local disk partition and a Network Block Device (NBD). A LifeKeeper supported file system can be mounted on a NetRAID device like any other storage device. In this case, the file system is called a replicated file system. LifeKeeper can protect both a NetRAID device and a replicated file system.

The NetRAID device can be created in a degraded mode, in which case data will be written only on the local disk partition. Extending the NetRAID device to another server will create the NBD device and make the network connection between the two servers. Data will be replicated in a synchronous mode after the NetRAID device has been extended and the initial synchronization between the servers has been completed.

The nbd-client process executes on the primary server and connects to the nbd-server process running on the backup server.



Resynchronization

The process by which LifeKeeper for Linux Data Replication performs resynchronization of mirrors is serial and not parallel. Therefore if you would like to replicate multiple disk partitions on a single server, the process will occur one at a time. The serial resynchronization is not a limitation of LifeKeeper for Linux Data Replication, but a limitation of the Linux Software RAID implementation.

Features

LifeKeeper for Linux Data Replication:

- Replicates data in synchronous mode to ensure data integrity.
- Integrates into the LifeKeeper Graphical User Interface for administration and monitoring.
- Fully supports LifeKeeper Application Recovery Kits, including Oracle, MySQL, DB2, Sendmail, Apache Web Server, Print Services, NFS Server, Informix, and Samba.
- Eliminates the need for an expensive external RAID.
- Offers robust error-handling capabilities.
- Automatically resynchronizes between the primary and backup server upon system recovery.
- Monitors the health of the underlying system components and performs a local recovery in the event of failure.
- Is fully functional over any Local Area Network or Wide Area Network.
- Provides replication that is completely transparent to the applications involved because it is done at the partition level.
- Supports the ext3 and reiserfs journaling filesystems (see the *Release Notes* for any kernel-specific patch information or restrictions regarding journaling filesystem support).
- Supports Stonith devices for I/O fencing. For details, refer to the white paper *Implementing STONITH Support in LifeKeeper for Linux Clusters* available at www.steeleye.com/support/documentation.

Failover using LifeKeeper for Linux Data Replication


Four examples of what happens during a failover using LifeKeeper for Linux Data Replication are provided below.

Consider the following:

Your LifeKeeper for Linux cluster consists of two servers, Server 1 and Server 2. The primary server is Server 1 and Server 2 is the backup server.

Scenario 1:

Server 1 has successfully completed its replication to Server 2 after which Server 1 becomes inoperable.



Result: Server 2 now takes on the role of primary server and Server 1 becomes the backup server. Server 2 will wait until Server 1 is again operational then initiate a full resynchronization from Server 2 to Server 1.


Scenario 2:

Considering scenario 1, Server 2 (still the primary server) becomes inoperable during the resynchronization with Server 1 (now the backup server).

Result: Because the resynchronization process did not complete successfully there is potential for data corruption. As a result, LifeKeeper will not attempt to failover the data replication resource to Server 1. Only when Server 2 becomes operable will LifeKeeper attempt to bring the data replication resource in-service (ISP) on Server 2.

Scenario 3:

Both Server 1 (primary) and Server 2 (target) become inoperable. Server 1 (primary) comes back up first.



Result: Server 1 will not bring the LKDR resource in-service. In this case you will need to delete the flag:

\$LKROOT/subsys/scsi/resources/netraid/\$TAG_data_corrupt

from Server 1 and manually bring the data replication resource into service (activate it) on Server 1.

Note: User must be certain that Server 1 was the last primary before removing the \$TAG_data_corrupt file. Otherwise data corruption might occur.

Scenario 4:

Both Server 1 (primary) and Server 2 (target) become inoperable. Server 2 (target) comes back up first.



Result: LifeKeeper will not bring the LKDR resource ISP on Server 2. When Server 1 comes back up, LifeKeeper will automatically bring the LKDR resource ISP on Server 1.

LifeKeeper for Linux Data Replication Requirements

Your LifeKeeper configuration should meet the following requirements **prior** to the installation of LifeKeeper for Linux Data Replication. See the *LifeKeeper for Linux Planning and Installation Guide* for instructions regarding the configuration of your LifeKeeper hardware and software.

Hardware Requirements

- **Servers** - Two LifeKeeper for Linux supported servers configured in accordance with the requirements described in the *LifeKeeper for Linux Planning and Installation Guide* and the *LifeKeeper Release Notes*, which are shipped with the product media.
- **IP Network Interface Cards** - Each server requires at least one Ethernet TCP/IP-supported network interface card. Remember, however, that a LifeKeeper cluster requires two communications paths; two separate LAN-based communications paths using dual independent sub-nets are recommended, and at least one of these should be configured as a private network. However using a combination of TCP and TTY is also supported.
Note: Due to the nature of software mirroring, network traffic between servers can be heavy. Therefore it is recommended that you implement a separate private network for your NetRAID devices, which may require additional network interface cards on each server.
- **Disk Partitions** - Disk partitions on the primary and backup server that will act as the source and target partition respectively. The target partition should be at least as large as the source partition.

Software Requirements

- **TCP/IP software** - Each server in your LifeKeeper configuration requires TCP/IP software.
- **LifeKeeper Software** - It is imperative that you install the same version of the LifeKeeper for Linux software and apply the same versions of the LifeKeeper for Linux software patches to each server in your LifeKeeper configuration.
- **LifeKeeper Installation Support** - In most cases, you will need to install the following package (see the “Product Requirements” section in the *LifeKeeper for Linux Release Notes* for specific LKDR requirements):

HADR-<Linux Distribution>-<Kernel Version>

This package must be installed on each server in your LifeKeeper cluster **prior** to the installation of LifeKeeper for Linux Data Replication. The HADR packages are located on the LifeKeeper Installation Support CD, and the appropriate package for your kernel version is automatically installed by the Installation Support **setup** script. This package installs SteelEye versions of the **nbd** and (for some kernel versions) **raid1** modules.

Note:

- Once installed, you can verify the correct modules by using **dmesg** or examining the */var/log/messages* file to ensure that “SteelEye” is included in the **nbd** and **raid1** (if installed) strings.
- If the automatic installation of the HADR rpm fails or if you do not see “SteelEye” in the **nbd** and **raid1** module strings, please refer to the “Setup Your LifeKeeper Environment

Using the LifeKeeper Installation Support CD” section in the *Planning and Installation Guide* for instructions on manually installing these files.

- **Data Replication software** - Each server in your LifeKeeper configuration requires LifeKeeper for Linux Data Replication software. LifeKeeper for Linux Data Replication is provided on a CD. It is packaged, installed and removed via the Red Hat Package Manager, rpm. The following rpm file is supplied on the LifeKeeper for Linux Data Replication CD: **steeleye-lkDR** - includes LifeKeeper for Linux Data Replication related files. Please see the *LifeKeeper for Linux Planning and Installation Guide* for specific instructions on the installation and removal of LifeKeeper for Linux Data Replication.

Configuring LifeKeeper for Linux Data Replication

This section describes LifeKeeper for Linux Data Replication configuration details, including the three data replication types. It also contains information you should consider before you start to configure and administer LifeKeeper for Linux Data Replication. Please refer to your *LifeKeeper Online Product Manual* for instructions on configuring LifeKeeper Core resource hierarchies.

LifeKeeper for Linux Data Replication Types

There are three different LifeKeeper for Linux Data Replication resource types:

Data Replication Resource will create/extend the NetRAID device and put it under LifeKeeper protection. You should format the NetRAID device (*dev/mdX*) and manually mount it on a mount point to use the file system outside of the LifeKeeper environment. You might choose this replication type if you want to format the NetRAID device with a file system that is not supported by LifeKeeper, or if you want to use it as a raw I/O device. Should you decide to protect the manually mounted file system with LifeKeeper, you can do so by creating/extending a file system hierarchy as if you were using a shared storage partition (assuming you are using a LifeKeeper supported file system type).

In order to allow the user continued data access, LifeKeeper for Linux Data Replication will not attempt to unmount and delete a NetRAID device if it is currently mounted. If there is a LifeKeeper protected filesystem on top of the NetRAID device, the file system recovery kit will be responsible for mounting/unmounting it during failovers. If, however, the user has created only a data replication resource with no LifeKeeper protected filesystem, the user has to handle the mounting/unmounting manually. This means that if the data replication resource is mounted, the user has to unmount it before attempting a manual failover and mount it on the other server after the manual failover.

New Replicated File System will create/extend the NetRAID device, mount the given mount point on the NetRAID device, and put both the LifeKeeper supported file system and the NetRAID device under LifeKeeper protection. The NetRAID device, along with the local partition will be formatted and all data will be deleted. **CAUTION: All data will be deleted.**

Replicating Existing File System will use a currently mounted partition and attempt to create a NetRAID device without deleting the data on the partition. Given a mount point, LifeKeeper for Linux Data Replication will check to see that the filesystem will fit on the NetRAID device. (Due to the NetRAID device chunk size, the resulting NetRAID device can be up to 4 kilobytes smaller than the original disk partition.) If this is true, it will unmount the mount point from the local partition, create the NetRAID device using the local partition and mount the mount point on the NetRAID device. It will then put both the NetRAID device and the LifeKeeper supported file system under LifeKeeper protection. If your filesystem will not fit on the NetRAID device that is to be created, then you may attempt either to resize the filesystem or, as a last resort, to use the **New Replicated File System** resource type. This may require you to backup your data from the local partition, and/or copy it to the NetRAID device after it has been created and formatted.

Specific Configuration Considerations

The following should be considered before operating LifeKeeper for Linux Data Replication:

- This release of LifeKeeper for Linux Data Replication does not support cascading failover or n-way mirroring (i.e. the LifeKeeper cluster is comprised of only two servers).
- This release of LifeKeeper for Linux Data Replication does not support standard LifeKeeper configurations using shared storage being replicated to an additional server.
- This release of LifeKeeper for Linux Data Replication does not support Automatic Switchback for data replication resources. Additionally, the Automatic Switchback restriction is applicable for any other LifeKeeper resource sitting on top of a LifeKeeper for Linux Data Replication resource.
- Users configuring data replication resources should avoid use of an interface/address already in use by a LifeKeeper IP resource with local recovery enabled. For example, if a LifeKeeper IP resource is configured on interface eth1 having local recovery enabled with interface eth2, data replication resources should avoid using either eth1 or eth2. Enabling local recovery will disable the interface during switchover to the backup interface, which can cause LifeKeeper for Linux Data Replication failure.
- This release of the LifeKeeper for Linux Data Replication supports SCSI, IDE, IBM ServeRAID (non-shared), and CPQARRAY disk environments.
- If you choose to create a [Data Replication Resource](#) (see the *Creating a Data Replication Resource* section), the NetRAID device will need to be formatted with a user preferred file system type and mounted on a local partition before it can be used. Once this has been accomplished, a file system resource can then be created and extended using the NetRAID device. Alternatively, you can simply use the NetRAID device as a raw I/O device.
- If you choose to create a [New File System Resource](#), your local partition will be formatted and all existing data on that file system will be deleted. **CAUTION: All data will be deleted.**
- If you choose the [Replicating an Existing File System](#) resource type, will check to see that the filesystem will fit on the NetRAID device. If your filesystem will not fit on the NetRAID device that is to be created, then you will need to create a new partition and a [New File System Resource](#).
- Keep in mind that with the use of each of the three data replication resource types, existing data on the [target partition](#) will be deleted and replaced with data from the source partition.
- Generally, during the mirroring of two partitions, the partition on the backup (target) server is locked to prevent write access by others. Unfortunately, the Linux operating system does not allow the partition on the backup server to be locked. Because Linux restricts LifeKeeper from locking your target partition, we strongly recommend that you avoid writing to this partition. Additionally, we strongly suggest that you refrain from selecting your unlocked target partition as a member of a different mirror. Please note that this locking issue is strictly a limitation of the Linux operating system and not of LifeKeeper for Linux Data Replication.
- When extending a Data Replication resource hierarchy from the primary server to the backup server, the size of the partition on the backup server must be equal to or greater than the size of the partition on the primary server.
- A LifeKeeper for Linux Data Replication resource (and its parent resources) cannot failover during the synchronization process.

- If your data replication resource is taken out of service (deactivated) during the synchronization of your primary and backup server, that resource can only be brought back into service (activated) on your primary server, and a full resynchronization will follow.
- If your primary server becomes inoperable during the synchronization process, your backup server will not be able to bring your data replication resource into service (activate it). Once your primary server becomes functional again, a full resynchronization of the servers will be initiated.
- If both your primary and backup server become inoperable, your data replication resource will be brought into service (activated) only when **both** servers are functional again. This is to avoid data corruption.
- If both servers become inoperable, when the primary server becomes functional again it will not bring the data replication resource hierarchy into service (activate it) until the backup server is also operational. This is to avoid data corruption that could result from initiating the resynchronization in the wrong direction. If you are certain that the only operable server was the last server on which the resource was “In Service Protected” (ISP), then delete the flag-file:
`$LKROOT/subsys/scsi/resources/netraid/$TAG_data_corrupt`
 and manually bring the data replication resource instance into service on this server.
- Although the NetRAID device replicates data in a synchronous mode, there is a very small interval during which data may not be synchronized. This is because LifeKeeper for Linux Data Replication monitors the health of the NetRAID device on a periodic preset interval. The default value is 120 seconds. This value is defined by the variable **LKCHECKINTERVAL** in `/etc/default/LifeKeeper`. See the **lkcheck(8)** man page for information on how to modify this value.

Configuration Tasks

You can perform all LifeKeeper for Linux Data Replication administrative tasks via the LifeKeeper Graphical User Interface (GUI). The LifeKeeper GUI provides a guided interface to configure, administer, and monitor LifeKeeper for Linux Data Replication resources.

Overview

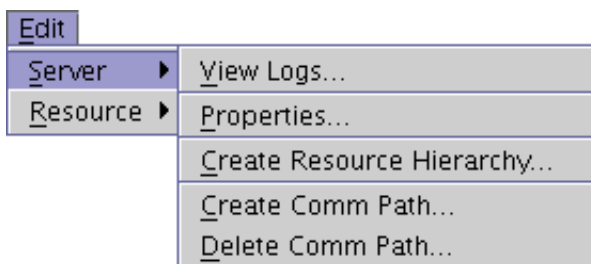
The following tasks are available for configuring LifeKeeper for Linux Data Replication:

- **Create a Resource Hierarchy** - Creates a Data Replication resource hierarchy
- **Delete a Resource Hierarchy** - Deletes a Data Replication resource hierarchy
- **Extend a Resource Hierarchy** - Extends a Data Replication resource hierarchy from the primary server to the backup server.
- **Unextend a Resource Hierarchy** - Unextends (removes) a Data Replication resource hierarchy from a single server in the LifeKeeper cluster.
- **Create Dependency** - Creates a child dependency between an existing resource hierarchy and another resource instance and propagates the dependency changes to all applicable servers in the cluster.
- **Delete Dependency** - Deletes a resource dependency and propagates the dependency changes to all applicable servers in the cluster.
- **In Service** - Activates a resource hierarchy.
- **Out of Service** - Deactivates a resource hierarchy.
- **View/Edit Properties** - View or edit the properties of a resource hierarchy.


Creating a Data Replication Resource Hierarchy

Perform the following on your primary server:

1. Select **Edit > Server > Create Resource Hierarchy**



2. The “**Select Recovery Kit**” dialog appears. Select the **Data Replication** option from the drop down list.

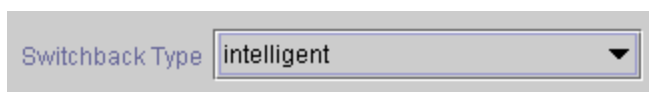


Click **Next** to continue.

CAUTION: If you click the **Cancel** button at any time during the sequence of creating your hierarchy, LifeKeeper will cancel the entire creation process.

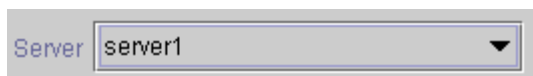
3. The “**Switchback Type**” dialog appears. The switchback type determines how the data replication resource will be switched back to the primary server when it becomes in-service (active) on the backup server following a failover. Switchback types are either *intelligent* or *automatic*. Intelligent switchback requires administrative intervention to switch the resource back to the primary server while automatic switchback occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.

CAUTION: This release of LifeKeeper for Linux Data Replication does not support Automatic Switchback for data replication resources. Additionally, the Automatic Switchback restriction is applicable for any other LifeKeeper resource sitting on top of a LifeKeeper for Linux Data Replication resource.



Click **Next** to continue.

4. The “**Server**” dialog appears. Select the name of the server where the NetRAID device will be created (typically this is your primary server). All servers in your cluster are included in the drop down list box.




Click **Next** to continue.

5. The “**Hierarchy Type**” dialog appears. Select the Data Replication type you wish to create by selecting one of the following:
 - **Data Replication Resource** - This option will only create a NetRAID device and place the device under LifeKeeper for Linux protection. You should select this option if you only want to create a data replication device on a disk partition and place the device under LifeKeeper protection. You will need to manually make and mount a file system on this device in order to create a readable mirror. You will need one free partition for this resource type.
 - **New Replicated File System** - This option will create a NetRAID device, format it with a LifeKeeper supported file system type, mount the file system on the NetRAID device, and place both the mounted file system and the NetRAID device under LifeKeeper protection. The NetRAID device and the local partition will be formatted, causing existing data to be deleted. You should select this option if you want to create a mirror on a new file system and place it under LifeKeeper protection. You will need one free partition for this resource type.

CAUTION: This option will cause your local partition to be formatted and all existing data will be deleted.
 - **Replicate Existing File System** - This option will take a currently mounted file system on a local partition, break the mount point on the local partition, create a NetRAID device, then re-mount the file system on the NetRAID device. Both the NetRAID device and the mounted file system are placed under LifeKeeper protection. You should select this option if you want to create a mirror on an existing file system and place it under

LifeKeeper protection. LifeKeeper for Linux Data Replication will check to see that the filesystem will fit on the NetRAID device. If it will not, you will need to select **New Replicated File System** as your **Hierarchy Type**.



Click **Next** to continue.

The next sequence of dialog boxes depends on which Hierarchy Type you have chosen. While some of the dialog boxes may be the same for each Hierarchy Type, their sequence and the required information may be slightly different. The next three sections take you through the remainder of the Hierarchy creation process.

Creating a Data Replication Resource

1. Once you have selected **Data Replication Resource** as your **Hierarchy Type**, the next dialog box will ask you to select the **Local Disk Partition** that will be replicated from the primary/template server to the backup/target server.

The list of Local Disk Partitions in the drop down box contains all the available partitions that are not:

- currently mounted
- swap type partitions
- LifeKeeper-protected partitions

The drop down list will also filter out special partitions, for example, root (`/`), boot (`/boot`), `/proc`, floppy and cdrom.



Click **Next** to continue.

2. Select or enter a unique **Data Replication Resource Tag** name for the Data Replication resource instance.



3. An information window appears notifying you that you will have to manually make the file system and mount the NetRAID device (`/dev/mdX`) before being able to use it.

You are about to create a new replicated device with /dev/sdf1 as the source partition.

You will have to manually make and mount the file system of your choice before being able to use this device.

Do you wish to continue?

Click **Continue** to create your data replication device on the local disk partition.

- An information box appears and LifeKeeper will verify that you have provided valid data to create your data replication resource. If LifeKeeper detects a problem, an ERROR will appear in the information box. If the validation is successful, your resource will be created.

Creating Data Replication Resource...

```
03/08/2001 11:15:34 create: BEGIN creation of "lkdr32335" on server "server1"
netraid_create: Creating the RAID1 (/dev/md0) device ...
netraid_create: /dev/sdf1 is configured to be mirrored through /dev/nb0
03/08/2001 11:15:36 create: END successful creation of "lkdr32335" on server "
```

Click **Next** to continue.

- An information box appears announcing the successful creation of your data replication resource device. You must **Extend** the hierarchy to another server in your cluster to begin data replication and in order to place it on the backup/target server and under LifeKeeper protection.

You have successfully created a resource hierarchy on one server. You may select continue in order to extend this resource hierarchy to another server, or you may cancel at this point.

If you cancel, the resource hierarchy provides no protection for your applications until it is extended to at least one other server in the cluster.

Click **Continue** to extend the resource.

Click **Cancel** if you wish to extend your resource at another time.

Verifying Integrity of Extended Hierarchy...

Hierarchy Verification Finished

WARNING: Your hierarchy exists on only one server. Your
 WARNING: application has no protection until you extend it
 WARNING: to at least one other server.

- Click **Done** to exit the Create Resource Hierarchy menu selection.

Creating a New Replicated File System

1. Once you have selected **New Replicated File System** as your **Hierarchy Type**, the next dialog box will ask you to select or enter the **Local Disk Partition** where the file system will be placed on the primary server.

The list of Local Disk Partitions in the drop down box contains all the available partitions that are not:

- already mounted
- swap type partitions
- LifeKeeper-protected partitions

The drop down list will also filter out special partitions, for example, root (/), boot (/boot), /proc, floppy and cdrom.

A screenshot of a dialog box with a label 'Local Disk Partition' and a dropdown menu. The dropdown menu is open and shows the selected option as '/dev/sdg1'.

Click **Next** to continue.

2. Select or enter the **New Mount Point** of the new file system. This should be the mount point where the replicated partition will be located.

A screenshot of a dialog box with a label 'New Mount Point' and a text input field. The input field contains the text '/test'.

Click **Next** to continue.

3. Select the **File System Type**. You may only choose from the LifeKeeper supported file system types.

A screenshot of a dialog box with a label 'New Filesystem Type' and a dropdown menu. The dropdown menu is open and shows the selected option as 'ext2'.

Click **Next** to continue.

4. Select or enter the **File System Resource Tag** name for the file system resource instance.

A screenshot of a dialog box with a label 'File System Resource Tag' and a text input field. The input field contains the text 'test-on-server1'.

5. Click **Create** for confirmation of the location where your new file system will be created and a warning indicating the reformat of your local partition. LifeKeeper's Create Resource Wizard will create your New Replicated File System resource hierarchy on the primary server.

```
WARNING: You are about to create a new filesystem with the following
attributes.
```

```
Disk Partition: /dev/sdg1
```

```
Mount Point: /test
```

```
Filesystem Type: ext2
```

```
Creating a new filesystem requires formatting /dev/sdg1.
```

```
ALL DATA ON /dev/sdg1 WILL BE LOST.
```

```
Do you wish to continue?
```

Click **Continue** to continue.

- LifeKeeper will verify that you have provided valid data to create your resource on a new file system. If LifeKeeper detects a problem, an **ERROR** will appear in the information box. If the validation is successful, your resource will be created.

Creating Data Replication Resource...

```
03/08/2001 11:22:51 create: BEGIN creation of "lkdr32716" on server "server1"
netraid_create: Creating the RAID1 (/dev/md1) device ...
netraid_create: /dev/sdg1 is configured to be mirrored through /dev/nb1
03/08/2001 11:22:53 create: END successful creation of "lkdr32716" on server1
```

CREATING RESOURCE HIERARCHIES

```
devicehier: Using /opt/LifeKeeper/lkadm/subsys/scsi/netraid/bin/devicehier to
Creating Resource Instance "test-on-server1" with id "/test" on machine "server1"
Resource "test-on-server1" Successfully Created on machine "server1"
Creating Dependency "test-on-server1"- "lkdr32716" on machine "server1":
Dependency "test-on-server1"- "lkdr32716" Successfully Created on machine "server1"
Removing /etc/fstab entry
```

Click **Next** to continue.

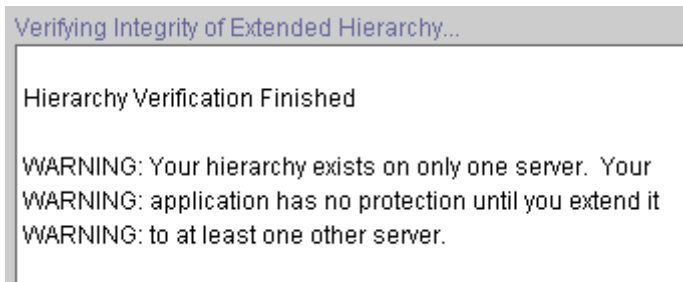
- An information box appears announcing the successful creation of your new replicated filesystem resource hierarchy. You must **Extend** the hierarchy to another server in your cluster to begin data replication and in order to place it under LifeKeeper protection.

```
You have successfully created a resource hierarchy on one server. You may
select continue in order to extend this resource hierarchy to another server,
or you may cancel at this point.
```

```
If you cancel, the resource hierarchy provides no protection for your
applications until it is extended to at least one other server in the cluster.
```

Click **Continue** to extend the resource.


Click **Cancel** if you wish to extend your resource at another time.



8. Click **Done** to exit the Create Resource Hierarchy menu selection.

Replicating an Existing File System

1. Once you have selected **Replicate Existing File System** as your **Hierarchy Type**, the next dialog box will ask you to select the **Existing Mount Point**. This should be the mount point to be mounted on the NetRAID device on the primary server. The local partition should already be mounted on this mount point and the filesystem should fit on the NetRAID device. If not, you will need to select **New Replicated File System** as your **Hierarchy Type**.



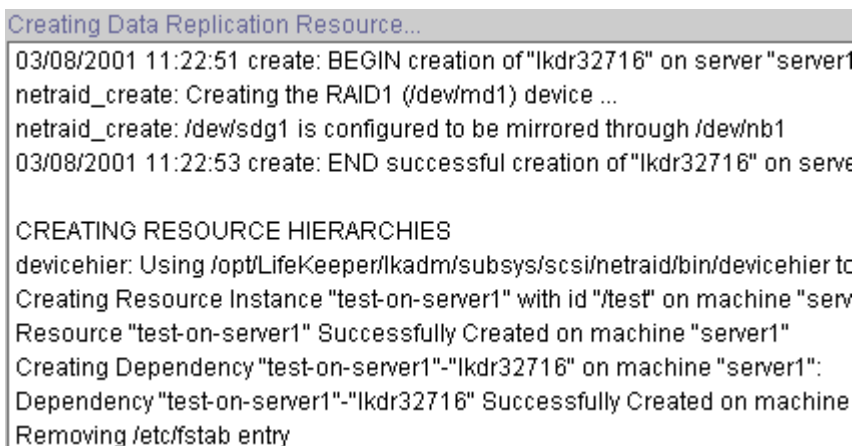
Click **Next** to continue.

2. Select or enter the **File System Resource Tag** name.



Click **Create** to create your data replication resource on the primary server.

3. LifeKeeper will verify that you have provided valid data to create your data replication resource. If LifeKeeper detects a problem, an **ERROR** will appear in the information box. If the validation is successful, your resource will be created.



Click **Next** to continue.

- An information box appears announcing that you have successfully created an existing replicated file system resource hierarchy. You must **Extend** the hierarchy to another server in your cluster to begin replication and to place it under LifeKeeper protection.

You have successfully created a resource hierarchy on one server. You may select continue in order to extend this resource hierarchy to another server, or you may cancel at this point.

If you cancel, the resource hierarchy provides no protection for your applications until it is extended to at least one other server in the cluster.

Click **Continue** to extend the resource.

Click **Cancel** if you wish to extend your resource at another time.

Verifying Integrity of Extended Hierarchy...

Hierarchy Verification Finished

WARNING: Your hierarchy exists on only one server. Your
 WARNING: application has no protection until you extend it
 WARNING: to at least one other server.

- Click **Done** to exit the Create Resource Hierarchy menu selection.

Deleting a Resource Hierarchy

To delete a data replication resource from all servers in your LifeKeeper configuration, complete the following steps:

- From the LifeKeeper GUI menu, select **Edit**, then **Resource**. From the drop down menu, select **Delete Resource Hierarchy**.
- Select the name of the **Target Server** where you will be deleting your data replication resource hierarchy.

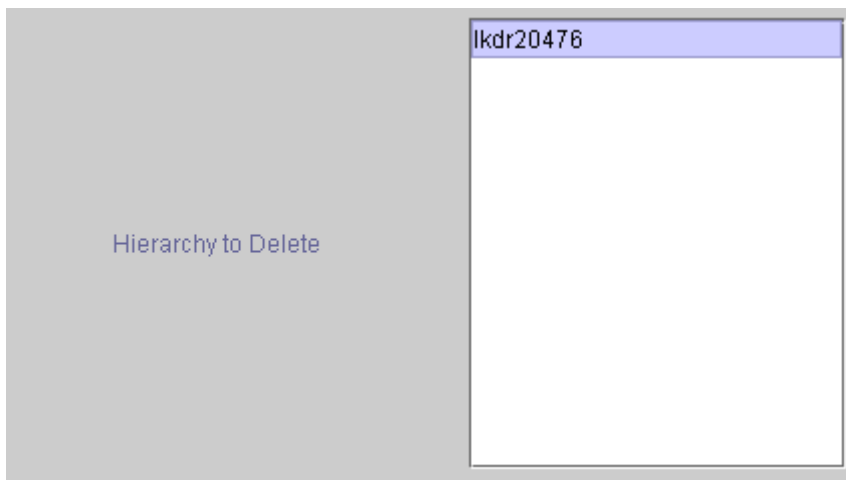
Note: If you selected the Delete Resource task by right-clicking from the right pane on an individual resource instance, or from the left pane on a global resource where the resource is on only one server, this dialog box will not appear.

Target Server

Click **Next** to continue.


- Select the **Hierarchy to Delete**. Identify the resource hierarchy you wish to delete, and highlight it.

Note: If you selected the Delete Resource task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to continue.

4. An information box appears confirming your selection of the target server and the hierarchy you have selected to delete.



Click **Delete** to continue.

5. An information box appears confirming that the data replication resource instance was deleted successfully.
6. Click **Done** to exit the Delete Resource Hierarchy menu selection.

Note: If the NetRAID device was mounted prior to the resource deletion then it will remain mounted. Otherwise, the NetRAID device will also be deleted.

Extending Your Hierarchy

After you have created a hierarchy, you should extend that hierarchy to another server in the cluster. There are three possible ways to extend your resource instance:

1. When you successfully create your data replication resource hierarchy you will have an opportunity to select **Continue**, which will allow you to proceed with extending your resource hierarchy to your backup server.
2. Right-click on an unextended hierarchy in either the left or right pane on the LifeKeeper GUI.
3. Select the “Extend Resource Hierarchy” task from the LifeKeeper GUI by selecting **Edit, Resource, Extend Resource Hierarchy** from the drop down menu. This sequence of selections will launch the Extend Resource Hierarchy wizard.


Note: The **Accept Defaults** button is intended for the user who is familiar with the LifeKeeper **Extend Resource Hierarchy** defaults and wants to extend a LifeKeeper resource

hierarchy with minimal input required. If you click **Accept Defaults**, you will be prompted only for the Local Disk Partition, Network Interface, and Data Replication Resource Tag.

Users who prefer to extend a LifeKeeper resource hierarchy using the interactive, step-by-step interface of the GUI dialogs should use the **Next** button.

- a. The first dialog box to appear will ask you to select the **Template Server** where your Data Replication resource hierarchy is currently in service. It is important to remember that the **Template Server** you select now and the **Tag to Extend** that you select in the next dialog box represent an in-service (activated) resource hierarchy. An error message will appear if you select a resource tag that is not in service on the template server you have selected. The drop down box in this dialog provides the names of all the servers in your cluster.

Note: If you are entering the Extend Resource Hierarchy task by continuing from the creation of a data replication resource hierarchy, this dialog box will not appear, because the wizard has already identified the template server in the create stage. This is also the case when you right-click on either the data replication resource icon in the left pane or right-click on the data replication resource box in the right pane of the GUI window and choose *Extend Resource Hierarchy*.




CAUTION: If you click the **Cancel** button at any time during the sequence of extending your hierarchy, LifeKeeper will cancel the extend hierarchy process. However, if you have already extended the resource to another server, that instance will continue to be in effect until you specifically unextend it.

Click **Next** to continue.


- b. Select the **Tag to Extend**. This is the name of the data replication instance you wish to extend from the template server to the target server. The wizard will list in the drop down list box all the resources that you have created on the template server.

Note: Once again, if you are entering the Extend Resource Hierarchy task immediately following the creation of an data replication resource hierarchy, this dialog box will not appear, since the wizard has already identified the tag name of your data replication resource in the create stage. This is also the case when you right-click on either the data replication resource icon in the left pane or on the data replication resource box in the right pane of the GUI window and choose *Extend Resource Hierarchy*.



Click **Next** to continue.

- c. Select the **Target Server** where you will extend your data replication resource hierarchy.



Click **Next** to continue.

- d. The **Switchback Type** dialog appears. The switchback type determines how the data replication resource will be switched back to the primary server when it becomes in service (active) on the backup server following a failover. Switchback types are either *intelligent* or *automatic*. Intelligent switchback requires administrative intervention to switch the resource back to the primary server while automatic switchback occurs as soon as the primary server is back on line and reestablishes LifeKeeper communication paths.

CAUTION: This release of LifeKeeper for Linux Data Replication does not support Automatic Switchback for data replication resources. Additionally, the Automatic Switchback restriction is applicable for any other LifeKeeper resource sitting atop a LifeKeeper for Linux Data Replication resource.

Click **Next** to continue.

- e. Select or enter a **Template Priority**. This is the priority for the Informix hierarchy on the server where it is currently in service. Any unused priority value from 1 to 999 is valid, where a lower number means a higher priority (1=highest). The extend process will reject any priority for this hierarchy that is already in use by another system. The default value is recommended. **Note:** This selection will appear only for the initial extend of the hierarchy.

Click **Next** to continue.

- f. Select or enter the **Target Priority**. This is the priority for the new extended data replication hierarchy relative to equivalent hierarchies on other servers. Any unused priority value from 1 to 999 is valid, indicating a server's priority in the cascading failover sequence for the resource. A lower number means a higher priority (1=highest). Note that LifeKeeper assigns the number "1" to the server on which the hierarchy is created by default. The priorities need not be consecutive, but no two servers can have the same priority for a given resource.

Click **Next** to continue.

- g. An information box appears explaining that LifeKeeper has successfully checked your environment and that all the requirements for extending this data replication resource have been met. If there are requirements that have not been met, LifeKeeper will disable the **Next** button, and enable the **Back** button.

```
Executing the pre-extend script...
Checking existence of extend and canextend scripts
Building independent resource list
Checking extendability for lkdr20476

Pre Extend checks were successful
```

Click **Back** to make changes to your resource extension.

Click **Cancel** to extend your resource another time.

Click **Next** to launch the Extend Resource Hierarchy configuration task.

The next sequence of dialog boxes depends on which Hierarchy Type you are extending. While some of the dialog boxes may be the same for each Hierarchy Type, their sequence and the required information may be slightly different. The next section lists the steps required to complete the extension of a data replication resource to another server. The following section lists the steps required to complete the extension of a new or existing replicated file system to another server.

Extending a Data Replication Resource

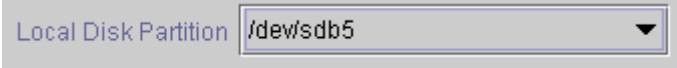
1. After you have been notified that your pre-extend script has executed successfully, the next dialog box to appear asks you to select the **Local Disk Partition** where the replicated file system will be located on the backup server.

The list of Local Disk Partitions in the drop down box contains all the available SCSI partitions that are not:

- already mounted
- swap type partitions
- LifeKeeper-protected partitions

The drop down list will also filter out special partitions, for example, root (`/`), boot (`/boot`), `/proc`, floppy and cdrom.

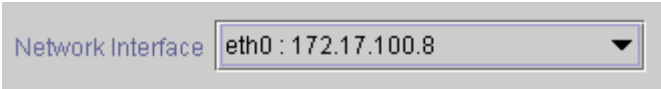
Note: The target partition must be greater than or equal to the source partition.



A screenshot of a configuration dialog box. On the left, the text 'Local Disk Partition' is displayed in a light blue font. To its right is a dropdown menu with a downward-pointing arrow. The selected option in the dropdown is '/dev/sdb5'.

Click **Next** to continue.

2. Select the **Network Interface** you will use for data replication traffic. Due to the nature of data replication, it is recommended you use a private (dedicated) network.



A screenshot of a configuration dialog box. On the left, the text 'Network Interface' is displayed in a light blue font. To its right is a dropdown menu with a downward-pointing arrow. The selected option in the dropdown is 'eth0 : 172.17.100.8'.

Click **Next** to continue.

3. Select or enter the **Data Replication Resource Tag** name.



A screenshot of a configuration dialog box. On the left, the text 'Data Replication Resource Tag' is displayed in a light blue font. To its right is a text input field containing the value 'lkdr26963'.

Click **Extend** to continue.

4. An information box will appear verifying that the extension is being performed.

```

Extending resource hierarchy...
Creating equivalencies
LifeKeeper Admin Lock (lkdr26963) Released

Hierarchy successfully extended
netraid_extend: Loading NBD module ...
netraid_extend: Loading RAID1 module ...
netraid_extend: Running the nbd-server on "server2" ...
netraid_extend: Running the nbd-client on "server1" ...
netraid_extend: Extending the RAID1 (/dev/md0) device from "server1" to
"server2"...
netraid_extend: "server1:/dev/sdb1" is being replicated to
"server2:/dev/sdb5"
    
```

LifeKeeper for Linux Data Replication supports only two server clusters - primary and backup. Therefore it cannot extend beyond one target node.

Click **Finish** to confirm the successful extension of your data replication resource instance.

5. Click **Done** to exit the Extend Resources Hierarchy menu selection.

Note: Be sure to test the functionality of the new instance on *both* servers.

Note: At this point LifeKeeper for Linux Data Replication has initiated the data resynchronization from the source to the target partition. The state of the data replication resource on the backup server is set to “Out of Service Failed” (OSF). Once the resynchronization is complete, the state will change to “Out of Service Unimpaired” (OSU).

During resynchronization, the data replication resource, and any resource that depends on it, will not be able to failover. This is to avoid data corruption.

Extending a New or Existing Replicated File System Resource

1. After you have been notified that your pre-extend script has executed successfully, the next dialog box to appear asks you to select the **Mount Point**. Enter the name of the file system mount point on the target server.

Click **Next** to proceed.

2. Select or enter the **Root Tag**. This is a unique name for the data replication resource instance on the target server.

Click **Next** to continue.

3. Select the **Local Disk Partition** where the replicated file system will be located on the backup server. This partition must be at least as large as the partition on the template server being replicated.

The list of Local Disk Partitions in the drop down box contains all the available SCSI partitions that are not:

- already mounted
- swap type partitions
- LifeKeeper-protected partitions

The drop down list will also filter out special partitions, for example, root (`/`), boot (`/boot`), `/proc`, floppy and cdrom.

Note: The target partition must be greater than or equal to the source partition.

Local Disk Partition

Click **Next** to continue.

4. Select the **Network Interface** you will use for data replication traffic. Due to the nature of data replication, it is recommended you use a private (dedicated) network.

Network Interface

Click **Next** to continue.

5. Select or enter the **Data Replication Resource Tag** name for the Data Replication Resource instance.

Data Replication Resource Tag

Click **Extend**.


6. An information box will appear verifying that the extension is being performed.

```
Extending resource hierarchy...
Creating equivalencies
LifeKeeper Admin Lock (lkdr26963) Released

Hierarchy successfully extended
netraid_extend: Loading NBD module ...
netraid_extend: Loading RAID1 module ...
netraid_extend: Running the nbd-server on "server2" ...
netraid_extend: Running the nbd-client on "server1" ...
netraid_extend: Extending the RAID1 (/dev/md0) device from "server1" to
"server2"...
netraid_extend: "server1:/dev/sdb1" is being replicated to
"server2:/dev/sdb5"
```

This release of LifeKeeper for Linux Data Replication supports only two server clusters - primary and backup. Therefore it cannot be extended beyond one target node.

Click **Finish** to confirm the successful extension of your resource instance.



7. Click **Done** to exit the Extend Resources Hierarchy menu selection.

Note: Be sure to test the functionality of the new instance on *both* servers.


Note: At this point LifeKeeper for Linux Data Replication has initiated the data resynchronization from the source to the target partition. The state of the data replication resource on the backup server is set to “Out of Service Failed” (OSF). Once the resynchronization is complete, the state will change to “Out of Service Unimpaired” (OSU).

During resynchronization, the data replication resource, and any resource that depends on it, will not be able to failover. This is to avoid data corruption.

Unextending Your Hierarchy

1. From the LifeKeeper GUI menu, select **Edit, Resource, and Unextend Resource Hierarchy**.
2. Select the **Target Server** where you want to unextend the data replication resource. It cannot be the server where the data replication resource is currently in service (active).


Note: If you selected the Unextend task by right-clicking from the right pane on an individual resource instance this dialog box will not appear.



Click **Next** to continue.


3. Select the Data Replication **Hierarchy to Unextend**.

Note: If you selected the Unextend task by right-clicking from either the left pane on a global resource or the right pane on an individual resource instance, this dialog will not appear.



Click **Next** to continue.

4. An information box appears confirming the target server and the data replication resource hierarchy you have chosen to unextend.



Click **Unextend**.

5. Another information box appears confirming that the data replication resource was unextended successfully.

6. Click **Done** to exit the Unextend Resource Hierarchy menu selection.

Note: At this point data is not being replicated to the backup server.

Testing Your Resource Hierarchy

You can test your data replication resource hierarchy by initiating a manual switchover. This will simulate a failover of the resource instance from the primary server to the backup server.

Performing a Manual Switchover from the LifeKeeper GUI

You can initiate a manual switchover from the LifeKeeper GUI by selecting **Edit, Resource, and In Service**. For example, an in-service request executed on a backup server causes the data replication resource hierarchy to be placed in-service on the backup server and taken out-of-service on the primary server. At this point, the original backup server is now the primary server and original primary server has now become the backup server.

Note: Manual failover is prevented for data replication resources during resynchronization.

If you execute the **Out of Service** request, the resource hierarchy is taken out-of-service without bringing it in-service on the other server. The resource can only be brought in-service on the same server, if it was taken out-of-service during resynchronization.

LifeKeeper GUI Display During Resynchronization

During the resynchronization of a data replication resource, the state of the data replication resource instance on the backup/target server is “Out of Service Failed” (OSF). However, the resource instance is “In Service Protected” (ISP) on the primary/template server. The LifeKeeper GUI reflects this status by showing the data replication resource on the backup server with a red “X” and the data replication resource on the primary server with a yellow triangle. Resources that depend on the data replication resource appear green. The global LifeKeeper status (displayed in the “Resource Hierarchy Tree” rectangle) appears as a yellow triangle. As soon as LifeKeeper detects that the resynchronization is complete, it sets the state of the data replication resource on the backup server to “Out of Service Unimpaired” (OSU), which is normal, and the GUI shows it as green on the primary server.

Monitoring the Status of Your NetRAID device

You can monitor the status of your LifeKeeper mirror by executing the following utility:

```
netraid_status <NetRAID Number>
```

Please reference the `netraid_status` manual pages for detailed information regarding the usage and functionality of this utility.

Additionally, you can determine the progress of your mirror as well as obtain general NetRAID device status information by monitoring the following file:

```
/proc/mdstat
```

Troubleshooting

If you have problems with your initial LKDR configuration, you should first ensure that the correct versions of the **nbd** and **raid1** modules have been installed if required by your kernel version. See the “Product Requirements” section in the *LifeKeeper for Linux Release Notes* for details about which kernels require the **nbd** and **raid1** modules. Use **dmesg** or examine the */var/log/messages* file. You should see the following:

```
nbd: (version SteelEye-n)
raid1: (version SteelEye-n)
```

where *n* is the specific module version that may vary. If you do not see “SteelEye” in the **nbd** and **raid1** strings, refer to the section entitled “Setup Your LifeKeeper Environment Using the LifeKeeper Installation Support CD” in the *Planning and Installation Guide* for instructions on manually installing these files.

Symptom	Possible Cause
NetRAID device not deleted after data replication resource deletion.	Deleting an LKDR resource will not delete the NetRAID device if the NetRAID device is mounted. You can manually unmount the device and delete it by executing: <i>\$LKROOT/kadm/subsys/scsi/netraid/bin/netraid_delete</i> (see NetRAID Utilities below for options).
Installation/HADR rpm fails	See the <i>Planning and Installation Guide</i> for complete instructions on manually installing these files.
Server Hung	Check your network connection. The primary server might appear hung while re-trying its writes to the backup server.
Failover Errors during failover	Check the status of your device. If resynchronization is in progress you cannot perform a failover.
After primary server panics, LKDR resource goes ISP on the secondary server, but when primary server reboots, the LKDR resource becomes OSF on both servers.	Check the “switchback type” selected when creating your data replication resource hierarchy. Automatic switchback is not supported for data replication resources in this release. You can change the Switchback type to “Intelligent” from the resource properties window.
Primary and Backup servers not synchronized	Check your “LKCHECKINTERVAL” setting in <i>/etc/default/LifeKeeper</i> . The value of this variable determines how often LifeKeeper for Linux Data Replication will monitor the health of your NetRAID device. Wait for this interval to elapse and check the status again.

<p>After both servers become inoperable, primary server cannot bring the resource ISP when it reboots.</p>	<p>If the primary server becomes operable before the secondary server, you will need to delete the file: <code>\$LKROOT/subsys/scsi/resources/netraid/\$TAG_data_corrupt</code> and manually bring the resource into service on the last known primary server.</p>
<p>Error creating a data replication hierarchy on currently mounted NFS file system</p>	<p>You are attempting to create a data replication hierarchy on a file system that is currently exported as NFS. You will need to replicate this file system before you export it.</p>
<p>Large number of messages appear in <code>/var/log/messages</code> log.</p>	<p>It is possible that a large number of messages similar to the following appear in <code>/var/log/messages</code>:</p> <p>kernel: mdX: write error while reconstructing, at block 0(1024)</p> <p>When a NetRAID device is constructed in degraded mode (no network connection yet) the Linux MD driver logs a large number of error messages in the system log when it retries to access the network part of the NetRAID device. This is normal behavior and the large number of retries is part of the Linux Kernel.</p> <p>Note: Due to the large number of error messages the size of the <code>/var/log/messages</code> file might become a problem if limited root partition disk space is an issue.</p>
<p>LKDR GUI wizard does not list a newly created partition</p>	<p>The Linux OS may not recognize a newly created partition until the next reboot of the system. View the <code>/proc/partitions</code> file for an entry of your newly created partition. If your new partition does not appear in the file, you will need to reboot your system.</p>
<p>LKDR fails to mount the NETRAID device during the creation of a New Replicated Filesystem</p>	<p>You may be using a newly created logical partition on a disk with only an extended partition (i.e. no primary partitions). In this case, you will need to manually execute the mkfs command on the newly created partition before attempting to create a New Replicated Filesystem.</p>

NetRAID Utilities

The NetRAID command line administration utilities allow the user to manually administer and maintain a NetRAID device. The following commands are provided for manual NetRAID administration and are located in the following location:

\$LKROOT/lkadm/subsys/scsi/netraid/bin

Note, however, that they can only be used in a LifeKeeper environment. All commands exit with a return value of 0 if they execute successfully. Please reference the NetRAID utilities manual pages for detailed information.

Utility	Usage
netraid_create	netraid_create <Template Partition> <NetRAID Number>
netraid_extend	netraid_extend <Template Server> <Template Partition> <Target Partition> <Target IP Address> <Port Number> <NetRAID Number>
netraid_delete	netraid_delete <Template partition> <NetRAID Number> <Target System> <Port Number>
netraid_unextend	netraid_unextend <Port Number>
netraid_resync	netraid_resync <Tag Name>
netraid_status	netraid_status <NetRAID Number> The possible return codes for netraid_status are: -3: Primary device has failed and backup device is not in sync with primary -2: Primary device has failed but backup device is functioning -1: NetRAID device <NetRAID Number> does not exist 0: Resynchronization has not begun (backup device is out of sync with primary device) 1-99: Resynchronization in progress (return code is percentage of resynchronization complete) 100: Resynchronization complete (backup device is in sync with primary device) WARNING: return codes -2 and -3 indicate that the primary disk device may be physically damaged and might need to be replaced – data loss may occur if the situation is not corrected immediately

Utility	Usage
netraid_info	<p>netraid_info <NetRAID Number></p> <p>Returns the following information:</p> <ul style="list-style-type: none"> NetRAID Device Name NetRAID Device Size Port Number used for NetRAID network connection NetRAID Status (Broken, Resync [%], Mirroring) <ul style="list-style-type: none"> Local Server Name [Local Server IP Address] Local Partition Name Local Partition ID Local LifeKeeper State (ISP, OSU, OSF, etc.) <ul style="list-style-type: none"> Remote Server Name [Remote Server IP Address] Remote Partition Name Remote Partition ID Remote LifeKeeper State (ISP, OSU, OSF, etc.)

Error Messages

Common Error Messages

Error Number	Error Message
000002	Usage error
000010	Error getting resource information
000011	Both Tag and ID name not specified
000019	Resource not found on local server
000022	END failed hierarchy <tag name> in service on server <servername>
000026	END failed ACTION for <tag name> on server <servername> due to <signal> signal

Hierarchy Creation

Error Number	Error Message
000012	Switchback type not specified
000013	Usage error
000014	Resource with either matching tag <tag name> or ID exists
000015	ins_create failed on server <server name>
000018	Error creating resource <tag name> on server <server name>
000021	Removing resource instance < tag name> from server <server name> due to an error during creation
000023	Error bringing resource < tag name> in service on server <server name>
000024	Failed resource creation of resource < tag name> on server <server name>
000027	Removing file system dependency from <parent tag> to <child tag> on server <server name> due to an error during creation
000028	Removing file system hierarchy <filesys tag> created by <parent tag> on server <server name> due to an error during creation
000029	<p>Switchback type mismatch between parent <parent tag> and child <child tag> on server <server name></p> <p>Action: Switchback mismatches can lead to unexpected behavior. You can manually alter switchback types for resources using the ins_setas command to eliminate this mismatch.</p>

000030	create: tag name not specified or extend: tag name not specified
--------	--

Hierarchy Extension

Error Number	Error Message
000003	Template resource < tag name> on server <server name> does not exist
000004	Template resource < tag name> cannot be extended to server <server name> because it already exists there
000005	Cannot access canextend script on server <server name>
000006	Cannot access extend script <path to extend> on server <server name>
000007	Cannot access depstoextend script <path to depstoextend> on server <server name>
000008	Cannot extend resource < tag name> to server <server name>
000009	Either <templatesys> or <templatetag> argument missing
000014	Resource with either matching tag < tag name> or ID exists
000015	ins_create failed on server <server name>
000018	Error creating resource < tag name> on server <server name>
000025	END failed resource extension of < tag name> on server <server name> due to a "<signal>" signal - backing out changes made to server
000030	create: tag name not specified or extend: tag name not specified

Restore

Error Number	Error Message
000023	Error bringing resource < tag name> in service on server <server name>

Resource Monitoring

Error Number	Error Message
000001	Calling sendevent for resource < tag name> on server <server name>

Data Replication Error Messages

Error Number	Error Message
104001	The "%s_data_corrupt" flag is set in "%s/subsys/scsi/resources/netraid/" on system "%s" Action: To avoid data corruption, LifeKeeper will not restore the resource. Remove the following flag: \$LKROOT/subsys/scsi/resources/netraid/\$TAG_data_corrupt
104002	Unable to open file "%s"
104003	Unable to start the resynchronization process for the NetRAID device "%s" of resource "%s" Action: Please check /proc/mdstat and the system log for error messages
104004	Unable to stop/remove the NetRAID device "%s" for resource "%s" on system "%s" Action: You can use the "raidstop(8)" command to manually stop the NetRAID device.
104005	Unable to get the NetRAID device number for resource tag "%s" on system "%s"
104006	Unable to get the NetRAID port number for resource tag "%s" on system "%s"
104007	Unable to get the local partition of the NetRAID device for resource tag "%s" on system "%s"
104008	Unable to determine the Local IP Address configured on network interface "%s" Action: Please check your network configuration
104009	Local Partition not available
104010	Unable to create the NetRAID device "%s" on system "%s"
104011	Resource "%s" is already extended to "%s"
104012	Unable to extend the NetRAID device "%s" to system "%s"
104013	The NetRAID device "%s" of resource "%s" is not in a mirroring state on system "%s" (Status : "%s") Action: You will not be able to failover the resource to another system. You will have to bring the resource in-service on this system first and a FULL resync will occur to ensure data reliability
104014	Resynchronization has begun for NetRAID device "%s" of resource "%s"
104015	The NetRAID device "%s" for resource "%s" has been successfully removed from system "%s"
104016	The NetRAID device "%s" has been successfully created on system "%s"

104017	The NetRAID device "%s" has been successfully extended to system "%s"
104018	Unable to Unextend the NetRAID device of resource "%s" from system "%s"
104019	The NetRAID device of resource "%s" has been successfully Unextended from system "%s"
104020	Cannot create the raidtab configuration file for NetRAID device "%s"
104021	Cannot kill the nbd-server listening on port "%s"
104022	Cannot load the "%s" kernel module Action: Please verify that this kernel module exists and that is not built-in to the kernel directly.
104023	Cannot create block device "%s"
104024	Cannot create the NetRAID device "%s" using the "mkraid" command
104025	Unextending resource "%s" while the NetRAID device "%s" is resynchronizing Action: The target partition might end up with corrupted data
104026	Deleting resource "%s" while the NetRAID device "%s" is resynchronizing Action: The target partition might end up with corrupted data
104027	Cannot get the remote system name for local tag "%s" Action: Please check the LifeKeeper resource equivalencies
104028	Cannot get the remote tag name for local tag "%s" Action: Please check the LifeKeeper resource equivalencies
104029	Cannot get the remote info list for local tag "%s" Action: Please check the LifeKeeper resource equivalencies
104030	Resynchronization has been completed for NetRAID device "%s" of resource "%s"
104031	Unable to determine the Remote IP Address of the system connected to the NBD-SERVER through port "%s"
104032	Unable to get the source partition size for resource tag "%s" on system "%s"
104033	Unable to set the NBD device size to "%s" for resource tag "%s" on system "%s"
104034	Another NBD client process is running on port "%s" Action: This process will be terminated
104040	Insufficient input parameters for "%s" creation

104041	Running on the wrong server, "%s"
104042	Specified existing Mount Point "%s" is not mounted
104043	Cannot get the block size of "%s" using the "\sfdisk -s\" command
104044	<p>To replicate an existing filesystem, the filesystem's size must be less than the size of the newly created NetRAID device. The existing filesystem is %s blocks(KB), but the newly created NetRAID device will be only %s blocks(KB)</p> <p>Action: You may be able to resize the existing filesystem, or you may need to archive the filesystem's data, create a new filesystem, and restore the data to the new filesystem</p>
104045	Mount Point "%s" is busy
104046	Cannot unmount existing Mount Point "%s"
104047	<p>Making the "%s" filesystem on "%s" ...</p> <p>This will erase all data on local ("%s") partition</p>
104048	Cannot make the %s filesystem on "%s"
104049	Cannot mount "%s" on "%s"
104050	Cannot create filesys hierarchy "%s"
104051	Invalid Data Replication resource type requested ("%s")
104052	Cannot get the hardware Id of device "%s"
104053	<p>The requested filesystem type "%s" is not supported by LifeKeeper</p> <p>Action: LifeKeeper supported filesystem types are listed under "\$LKROOT/lkadm/subsys/gen/filesys/supported_fs"</p>
104054	The following command failed to execute properly: %s

LifeKeeper GUI Related Errors

Error Number	Error Message
104901	The mount point %s is mounted Action: Please specify a mount point that is not mounted
104902	The mount point %s is not an absolute path Action: Please specify a mount point that begins with a slash
104903	The mount point %s is not empty Action: Please specify a mount point that does not exist or is empty
104904	Killing the nbd-client process has timed out. Action: This is not user correctable. Most probably the daemon is stuck in the kernel waiting for an event. LifeKeeper will retry the recovery after LKCHECKINTERVAL seconds
104905	Failed to set the LKDR resynchronization speed limit in %s