



LifeKeeper[®] for Linux v5.1
DRBD Recovery Kit
Administration Guide

March 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 © 2005 - 2006
By SteelEye Technology, Inc.
Palo Alto, CA U.S.A.
All Rights Reserved

Table of Contents

Introduction	3
Document Contents	3
Documentation and References	3
Requirements	5
Hardware Requirements	5
Software Requirements	5
Notes and Restrictions	6
Network Configuration	6
WAN Configuration	6
Restrictions	6
DRBD Administration	7
Using DRBD with LifeKeeper	7
DRBD Resource Hierarchy	7
Hierarchy Creation	7
Viewing DRBD Device Status	8
Disconnecting and Reconnecting the DRBD Replication Path	9
Troubleshooting	10

DRBD Recovery Kit Administration Guide

Introduction

The LifeKeeper for Linux DRBD Recovery Kit enables the use of the DRBD (Distributed Remote Block Device) data replication software in LifeKeeper environments. This allows LifeKeeper-protected applications to operate in non-shared storage configurations.

The DRBD Recovery Kit is different from most other LifeKeeper recovery kits in that it is never used alone; DRBD resources are always dependencies of other LifeKeeper resources. As such, many of the operations typically associated with a LifeKeeper recovery kit – for example, creating a resource hierarchy – are not directly applicable to the DRBD Recovery Kit.

Document Contents

This guide contains the following topics:

- [Documentation and References](#). Provides a list of related documentation along with the location where each document can be found.
- [Requirements](#). Lists the hardware and software necessary to properly setup, install, and operate the LifeKeeper for Linux DRBD Recovery Kit.
- [Notes and Restrictions](#). Lists some important notes and restrictions that apply to the DRBD Recovery Kit.
- [DRBD Administration](#). Provides information about using DRBD under LifeKeeper protection.
- [Troubleshooting](#). Lists some problems (along with potential solutions) that may be encountered while using the DRBD Recovery Kit.

Documentation and References

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 other LifeKeeper Recovery Kits, is available on the SteelEye Technology, Inc. website at:

www.steeleye.com/support/documentation

The LifeKeeper for Linux DRBD Recovery Kit protects DRBD devices. The kit does not configure or administer DRBD devices. Any support required for these tasks and/or any issues encountered with DRBD itself should be addressed by the provider of DRBD or by referring to the DRBD website, www.drbd.org.

The *DRBD HOWTO*, which is available at <http://www.slackworks.com/~dkrovich/DRBD/>, contains useful information about acquiring, building, and configuring the DRBD software.

Requirements

Your LifeKeeper configuration should meet the following requirements **prior** to the installation of the LifeKeeper for Linux DRBD Recovery Kit. Please see the *LifeKeeper for Linux Planning and Installation Guide* for specific instructions regarding the configuration of your LifeKeeper hardware and software.

Hardware Requirements

- **Servers** - Two LifeKeeper for Linux supported servers.
- **Network Interface Cards** - 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 DRBD 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.

Software Requirements

- **Operating System** - DRBD can be used with any major Linux distribution. See the *LifeKeeper for Linux Release Notes* for a list of supported distributions.
- **DRBD Software** - The DRBD software must be installed and configured prior to use of the DRBD Recovery Kit. See the *LifeKeeper for Linux Release Notes* for a list of supported DRBD versions.
- **LifeKeeper Software** - You must install the same version of the LifeKeeper Core on each of your servers. You must also install the same version of each recovery kit that you plan to use on each server. See the *LifeKeeper for Linux Release Notes* for specific LifeKeeper requirements.
- **LifeKeeper for Linux DRBD Recovery Kit** - The DRBD Recovery Kit is provided on a CD. It is packaged, installed, and removed with the Red Hat Package Manager, rpm. The rpm file containing the DRBD Recovery Kit is: **steeleye-lkDRBD**.

Refer to the *LifeKeeper for Linux Planning and Installation Guide* for instructions on installing or removing the LifeKeeper Core and the DRBD Recovery Kit.

The DRBD Recovery Kit must be installed on each server where DRBD is being used to replicate data that will be protected by LifeKeeper.

Notes and Restrictions

The following notes and restrictions apply to the DRBD Recovery Kit.

Network Configuration

When configuring DRBD devices, avoid using an interface/address already in use by a LifeKeeper IP resource that has local recovery enabled. For example, if a LifeKeeper IP resource is configured on interface eth1 having local recovery enabled with interface eth2, DRBD resources should avoid using either eth1 or eth2. Enabling local recovery will disable the interface during switchover to the backup interface, which can cause DRBD to fail.

WAN Configuration

Using LifeKeeper in a WAN environment requires special configuration due to the nature of WAN networking. The following tips are recommended:

- To prevent false failover, you should enable manual failover confirmation. (Edit the file `/etc/default/LifeKeeper` to set `CONFIRMSODEF=1` and create the LifeKeeper flag `confirmso!uname` using the `flg_create(8)` command). Because most WANs are somewhat less reliable than LANs and because typical WAN configurations will have only one comm path, this is usually a good idea. With this option enabled, a LifeKeeper failover will proceed only if the user confirms the failover by using the `lk_confirmso(8)` command. Refer to the `lk_confirmso` man page for more details.
- If the WAN link experiences periods of downtime in excess of 15 seconds on a regular basis, it may also be wise to tune the LifeKeeper heartbeat parameters. See “Tuning the LifeKeeper Heartbeat” in the *LifeKeeper Online Product Manual* for details.

Restrictions

The current release of the LifeKeeper for Linux DRBD Recovery Kit does not support the following:

- Cascading Failover
- Automatic Switchback option
 - All resource hierarchies that contain a DRBD resource should be configured for Intelligent Switchback.
- Replication of LifeKeeper-protected shared storage to an additional (third) server

DRBD Administration

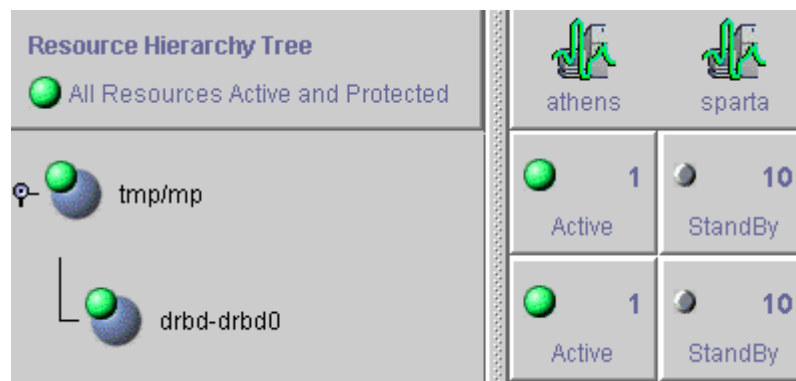
Using DRBD with LifeKeeper

DRBD can be used to replicate data (either synchronously or asynchronously) between two servers in a LifeKeeper cluster as an alternative to using shared storage. The LifeKeeper DRBD Recovery Kit provides the support for other LifeKeeper recovery kits to operate properly on top of DRBD devices.

DRBD Resource Hierarchy

The creation, extension, un-extension, and deletion of DRBD resources are always driven by the corresponding hierarchy administration processes of a higher-level resource.

The following example shows a typical DRBD resource hierarchy as it appears in the LifeKeeper GUI:



The resource *drbd-drbd0* is a DRBD resource, and the parent resource *tmp/mp* is a filesystem resource. The *drbd-drbd0* resource corresponds to the *drbd0* DRBD resource defined in the */etc/drbd.conf* file. The filesystem */tmp/mp* is mounted on the drbd device, which is currently replicating data from the server *athens* to the server *sparta*.

Hierarchy Creation

DRBD resources are created automatically during the hierarchy creation process of resources that sit on top of DRBD devices.

In order to create a hierarchy in which a higher-level application uses a DRBD device, the following procedure should be followed. (Please refer to the *DRBD HOWTO*, which is referenced in the **Documentation and References** section of this document, if you need more information about installing and configuring the DRBD software.)

1. Obtain a copy of a supported version of the DRBD software. Your Linux distribution may already be installed with a supported version of DRBD, in which case, skip to step 3.
2. Build and install the DRBD software on your systems.

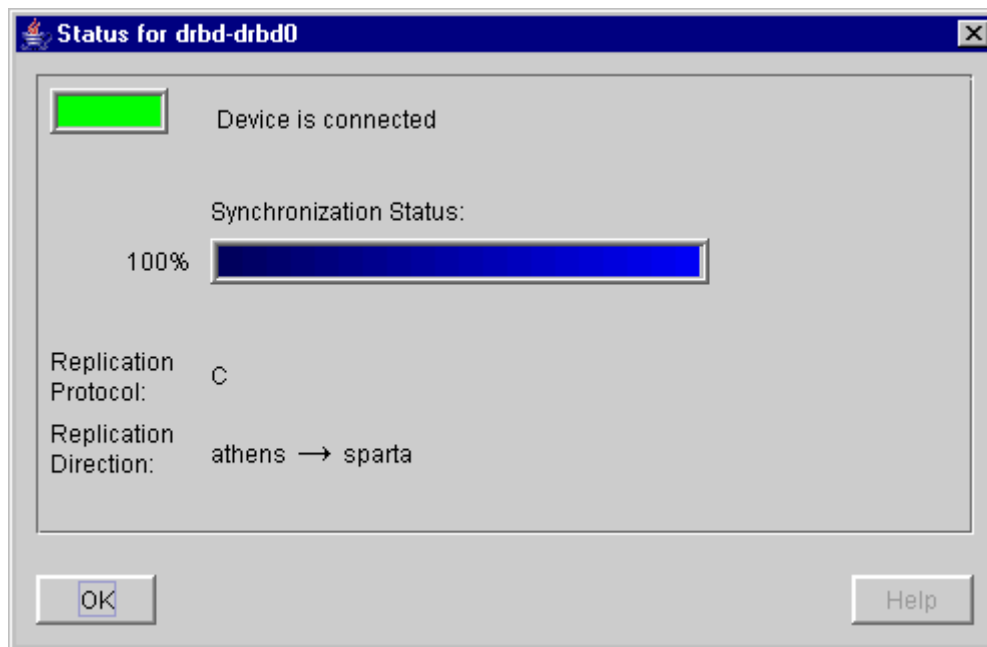
3. Make sure the drbd rc script (`/etc/init.d/drbd`) is configured to be started in runlevels 3 and 5 (at least). LifeKeeper will not be able to activate the DRBD devices if the rc script has not been started.
4. Configure a DRBD device to replicate between two systems in the LifeKeeper cluster (make sure that the `/etc/drbd.conf` file is identical on the two servers).
5. Create a filesystem on the newly created DRBD device, e.g., `mke2fs -j /dev/drbd0`, and mount the filesystem. Alternatively, you may bind a raw I/O device to the DRBD device.
6. Configure the application on top of the newly created filesystem, following the instructions in the administration guide for the recovery kit associated with the application.
7. Create and extend the application resource hierarchy following the instructions in the administration guide for the appropriate application recovery kit.

Viewing DRBD Device Status

Once you have created a resource hierarchy containing a DRBD device, you can view the status of the DRBD device using the LifeKeeper GUI.

To view the status of a DRBD device, do the following:

1. Right-click the DRBD resource in the LifeKeeper GUI.
2. From the pop-up menu, select **Properties**.
3. Click the **Status** button at the bottom of the **Resource Properties** dialog.



Disconnecting and Reconnecting the DRBD Replication Path

The DRBD replication path can be disconnected and reconnected through the LifeKeeper GUI.

Disconnecting the replication path will stop all writes from being replicated to the secondary system. You may wish to disconnect the replication path, for example, in order to make a snapshot of the secondary device or to increase write performance during peak traffic times.

To disconnect the replication path of a DRBD device that is in the *Connected* state, do the following:

1. Right-click the DRBD resource in the LifeKeeper GUI.
2. From the pop-up menu, select **Properties**.
3. Click the **Actions** button at the bottom of the **Resource Properties** dialog.
4. The default action in the dropdown menu is **Disconnect**:



Click **Next**.

5. You will be asked if you wish to continue. Click **Continue** to confirm.

To reconnect the DRBD replication path following a manual disconnect action or in cases (such as a “split brain” scenario) where DRBD is unable to reconnect automatically, do the following:

1. Right-click the DRBD resource in the LifeKeeper GUI.
2. From the pop-up menu, select **Properties**.
3. Click the **Actions** button at the bottom of the **Resource Properties** dialog.
4. The default action in the dropdown menu is **Reconnect**:



Click **Next**.

5. You will be asked if you wish to continue. Click **Continue** to confirm.

Troubleshooting

The following table lists some problems you may encounter, along with some possible solutions for each problem.

Problem	Suggested Solution
Resources appear green (ISP) on both primary and backup servers	<p>This is a “split brain” scenario caused by a temporary communication failure in the LifeKeeper cluster. After communication has been restored, both systems think they are the primary.</p> <p>DRBD will not reconnect when the devices on both systems are marked as “primary”. Manual intervention is required.</p> <p>You will need to determine which of the two resources was the secondary last. You should then take that resource out of service, leaving the original primary resource in service. You will then need to open the Resource Properties dialog in the GUI, click Actions, and select Reconnect. Click Continue to confirm, then Finish. Verify that the DRBD connection has been re-established by clicking the Status button in the Resource Properties dialog.</p>
<p>Resource fails to come in service on secondary system after the primary fails</p> <p><i>OR</i></p> <p>Resource fails to come in service on secondary system during manual switchover</p>	<p>The secondary leg of the DRBD device is in an inconsistent state (meaning a resync was in progress when the primary failed). You must bring the device back into service on the primary system and allow the resync to finish before the device can be brought into service on the secondary system.</p>