

Novell SUSE® Linux Enterprise Server

10

August 15, 2006

STORAGE ADMINISTRATION GUIDE
FOR EVMS

www.novell.com



Novell®

Legal Notices

Novell, Inc., makes no representations or warranties with respect to the contents or use of this documentation, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to revise this publication and to make changes to its content, at any time, without obligation to notify any person or entity of such revisions or changes.

Further, Novell, Inc., makes no representations or warranties with respect to any software, and specifically disclaims any express or implied warranties of merchantability or fitness for any particular purpose. Further, Novell, Inc., reserves the right to make changes to any and all parts of Novell software, at any time, without any obligation to notify any person or entity of such changes.

Any products or technical information provided under this Agreement may be subject to U.S. export controls and the trade laws of other countries. You agree to comply with all export control regulations and to obtain any required licenses or classification to export, re-export or import deliverables. You agree not to export or re-export to entities on the current U.S. export exclusion lists or to any embargoed or terrorist countries as specified in the U.S. export laws. You agree to not use deliverables for prohibited nuclear, missile, or chemical biological weaponry end uses. Please refer to www.novell.com/info/exports/ for more information on exporting Novell software. Novell assumes no responsibility for your failure to obtain any necessary export approvals.

Copyright © 2006 Novell, Inc. All rights reserved. No part of this publication may be reproduced, photocopied, stored on a retrieval system, or transmitted without the express written consent of the publisher.

Novell, Inc., has intellectual property rights relating to technology embodied in the product that is described in this document. In particular, and without limitation, these intellectual property rights may include one or more of the U.S. patents listed at <http://www.novell.com/company/legal/patents/> and one or more additional patents or pending patent applications in the U.S. and in other countries.

Novell, Inc.
404 Wyman Street, Suite 500
Waltham, MA 02451
U.S.A.
www.novell.com

Online Documentation: To access the online documentation for this and other Novell products, and to get updates, see www.novell.com/documentation.

Novell Trademarks

For a list of Novell trademarks, see the [Novell Trademark and Service Mark list \(http://www.novell.com/company/legal/trademarks/tmlist.html\)](http://www.novell.com/company/legal/trademarks/tmlist.html).

Third-Party Materials

All third-party trademarks and copyrights are the property of their respective owners.

Some content in this document is copied, distributed, and/or modified from the following document under the terms specified in the document's license.

EVMS User Guide, January 18, 2005

Copyright © 2005 IBM

License Information

This document may be reproduced or distributed in any form without prior permission provided the copyright notice is retained on all copies. Modified versions of this document may be freely distributed provided that they are clearly identified as such, and this copyright is included intact.

Contents

About This Guide	9
1 Overview of EVMS	11
1.1 Benefits of EVMS	11
1.2 Plug-In Layers	11
1.3 File Systems Support	12
1.4 Terminology	12
1.5 Mount Location of Storage Objects with EVMS	13
2 Using EVMS to Manage Devices	15
2.1 Configuring the System Device at Install to Use EVMS	15
2.1.1 Before the Install	15
2.1.2 During the Install	17
2.1.3 After the Install	20
2.2 Configuring the System Device After the Install to Use EVMS	22
2.2.1 Disable the boot.lvm and boot.md Services	23
2.2.2 Enable the boot.evms Service	23
2.2.3 Edit the /etc/init.d/boot.evms Script	23
2.2.4 Edit the /etc/fstab File	24
2.2.5 Edit the Boot Loader File	25
2.2.6 Force the RAM Disk to Recognize the Root Partition	26
2.2.7 Reboot the Server	27
2.2.8 Verify that EVMS Manages the Boot, Swap, and Root Partitions	27
2.3 Configuring LVM Devices After Install to Use EVMS	28
2.4 Using EVMS with iSCSI Volumes	28
2.5 Using the ELILO Loader Files (IA-64)	29
2.6 Starting EVMS	29
2.7 Starting the EVMS UI	29
3 Loading EVMS Devices by UUIDs	31
3.1 Naming Devices with udev	31
3.2 Understanding UUIDs	31
3.2.1 Using UUIDs to Assemble or Activate Devices	31
3.2.2 Finding the UUID for a Device	32
3.3 Using UUIDs in the Boot Loader and /etc/fstab File (x86)	32
3.4 Using UUIDs in the Boot Loader and /etc/fstab File (IA64)	33
4 Managing Devices	35
4.1 Understanding Disk Segmentation	35
4.1.1 Segment Managers	35
4.1.2 Disk Segments	36
4.2 Initializing Disks	36
4.2.1 Before You Begin	36
4.2.2 Guidelines	37
4.2.3 Adding a Segment Manager	37
4.3 Removing the Segment Manager from a Device	38

4.4	Creating Disk Segments (or Partitions)	38
4.5	Configuring Mount Options for Devices	39
4.6	What's Next	41
5	Managing Multipathing for Devices and Software RAIDs	43
5.1	Understanding Multipathing	43
5.1.1	What Is Multipathing?	43
5.1.2	Benefits of Multipathing	44
5.1.3	Guidelines for Multipathing	44
5.1.4	Device Mapper Multipath I/O Module	45
5.1.5	Multipath Tools	46
5.2	Before You Begin	46
5.2.1	Preparing the SAN Devices	46
5.2.2	Configuring mdadm.conf and lvm.conf for Multipathing	47
5.2.3	Configuring Device Drivers in initrd	47
5.3	Adding multipathd to the Boot Sequence	48
5.3.1	YaST	48
5.3.2	Command Line	48
5.4	Starting Multipath I/O Services	48
5.5	Configuring Multipath Settings in the /etc/multipath.conf File	48
5.6	Resolving Stalled I/O	49
5.7	Configuring Time-Out Settings for the HBA	49
5.8	Configuring Multipath I/O for Devices	49
5.9	Configuring Multipathing for an Existing Software RAID	51
5.10	Adding Storage Dynamically for Use with Multipath I/O	53
5.11	What's Next	54
6	Managing Software RAIDs	55
6.1	Understanding Software RAIDs on Linux	55
6.1.1	What Is a Software RAID?	55
6.1.2	Overview of RAID Levels	56
6.1.3	Comparison of RAID Performance	57
6.1.4	Comparison of Disk Fault Tolerance	58
6.1.5	Configuration Options for RAIDs	58
6.1.6	Guidelines for Member Disks	58
6.1.7	RAID 5 Algorithms for Distributing Stripes and Parity	59
6.1.8	Device Mapper	61
6.1.9	EVMS and the Multi-Disk (md) Driver	61
6.2	Creating and Configuring a Software RAID	61
6.3	Expanding a RAID	65
6.3.1	Adding Mirrors to a RAID 1 Device	65
6.3.2	Adding Segments to a RAID 4 or 5	66
6.4	Adding or Removing a Spare Disk	66
6.4.1	Do I Need a Spare Disk?	66
6.4.2	Adding a Spare Disk When You Create the RAID	66
6.4.3	Adding a Spare Disk to an Existing RAID	66
6.4.4	Removing a Spare Disk from a RAID	67
6.5	Removing a Failed Disk	67
6.6	Replacing a Failed Disk	68
6.7	Monitoring Status for a RAID	68
6.7.1	Monitoring Status with EVMSGUI	69
6.7.2	Monitoring Status with /proc/mdstat	69
6.7.3	Monitoring Status with mdadm	70

6.7.4	Monitoring a Remirror or Reconstruction.	71
6.8	Deleting a Software RAID and Its Data	72
7	Installing and Managing DRBD Services	73
7.1	Understanding DRBD	73
7.2	Installing DRDB Services.	73
7.3	Configuring the DRDB Service	74
8	Troubleshooting EVMS Devices, RAIDs, and Volumes	77
8.1	EVMS Volumes Might Not Appear When Using iSCSI	77

About This Guide

This guide provides information about how to manage storage devices on a SUSE® Linux Enterprise Server 10 server with the Enterprise Volume Management System (EVMS) 2.5.5 storage management tool.

- Chapter 1, “Overview of EVMS,” on page 11
- Chapter 2, “Using EVMS to Manage Devices,” on page 15
- Chapter 3, “Loading EVMS Devices by UUIDs,” on page 31
- Chapter 4, “Managing Devices,” on page 35
- Chapter 5, “Managing Multipathing for Devices and Software RAIDs,” on page 43
- Chapter 6, “Managing Software RAIDs,” on page 55
- Chapter 7, “Installing and Managing DRBD Services,” on page 73
- Chapter 8, “Troubleshooting EVMS Devices, RAIDs, and Volumes,” on page 77

Audience

This guide is intended for system administrators.

Feedback

We want to hear your comments and suggestions about this manual and the other documentation included with this product. Please use the User Comments feature at the bottom of each page of the online documentation, or go to www.novell.com/documentation/feedback.html and enter your comments there.

Documentation Updates

For the most recent version of the *SUSE Linux Enterprise Server 10 Storage Administration Guide for EVMS*, visit the [Novell Documentation Web site for SUSE Linux Enterprise Server 10 \(http://www.novell.com/documentation/sles10\)](http://www.novell.com/documentation/sles10).

Additional Documentation

For information about managing storage with the Linux Volume Manager (LVM), see the *SUSE Linux Enterprise Server 10 Installation and Administration Guide* (<http://www.novell.com/documentation/sles10>).

Documentation Conventions

In Novell documentation, a greater-than symbol (>) is used to separate actions within a step and items in a cross-reference path.

A trademark symbol (®, ™, etc.) denotes a Novell trademark. An asterisk (*) denotes a third-party trademark.

Overview of EVMS

1

The Enterprise Volume Management System (EVMS) 2.5.5 management tool for Linux* is an extensible storage management tool that integrates all aspects of volume management, such as disk partitioning, the Logical Volume Manager (LVM), the Multiple-Disk (MD) manager for software RAIDs, the Device Mapper (DM) for multipath I/O configuration, and file system operations.

- [Section 1.1, “Benefits of EVMS,” on page 11](#)
- [Section 1.2, “Plug-In Layers,” on page 11](#)
- [Section 1.3, “File Systems Support,” on page 12](#)
- [Section 1.4, “Terminology,” on page 12](#)
- [Section 1.5, “Mount Location of Storage Objects with EVMS,” on page 13](#)

1.1 Benefits of EVMS

EVMS provides the following benefits:

- Is an open source volume manager
- Provides a plug-in framework for flexible extensibility and customization
- Allows plug-ins to extend functionality for new or evolving storage managers
- Supports foreign partition formats
- Is cluster-aware

1.2 Plug-In Layers

EVMS abstracts the storage objects in functional layers to make storage management more user-friendly. The following table describes the current EVMS plug-in layers for managing storage devices and file systems:

Table 1-1 EVMS Plug-In Layers

Storage Managers	Description	Plug-Ins
Device	Manages the physical and logical disks	Device Mapper (DM)
Segment	Manages the partitioning of physical and logical disks into smaller segments of free space	DOS, GPT, System/390* (S/390), Cluster, BSD, Mac, and BBR For more information, see Section 4.1, “Understanding Disk Segmentation,” on page 35 .
Regions	Manages the combination of multiple storage objects	LVM/LVM2 for containers and region, MD for RAIDs, and DM for multipath I/O

Storage Managers	Description	Plug-Ins
EVMS Features	Manages EVMS features	Drive linking (linear concatenation), Bad Block Relocation (BBR), and Snapshot
File System Interface Modules (FSIM)	Manages the interface between the file system managers and the segment managers	For information, see Section 1.3, "File Systems Support," on page 12.
Cluster Manager Interface Modules	Manages the interface between the cluster manager and the file systems and devices	HeartBeat 2

1.3 File Systems Support

EVMS supports the following Linux file systems:

- ReiserFS (the default file system for SUSE® Linux Enterprise Server 10)
- EXT3
- XFS
- OCFS2
- JFS
- EXT2
- Swap
- NTFS (read only)
- FAT (read only)

For more information about file systems supported in SUSE Linux Enterprise Server 10, see the *SUSE Linux Enterprise Server 10 Installation and Administration Guide*. (<http://www.novell.com/documentation/sles10>).

1.4 Terminology

EVMS uses the following terminology in the EVMS user interface:

Table 1-2 *EVMS Terms*

Term	Description
Sector	The lowest level that can be addressed on a block device.
Disk	The lowest level of a device. This can be a physical disk or physically ordered contiguous sectors that form a logical disk.
Segment	Partitions and other memory regions on the disk such as the master boot record (MBR). The underlying mapping can be to disks or other segments.

Term	Description
Region	An ordered set of logically contiguous sectors that might or might not be physically contiguous. The underlying mapping can be to logical disks, disk segments, or other storage regions. For example, the “volumes” exported by the DM manager for multipath I/O are considered regions in the EVMS architecture.
Feature (Feature Object, EVMS Feature, EVMS Object)	A logically contiguous address space created from one or more disks, segments, regions, or other feature objects through the use of an EVMS feature.
Storage Object	Any memory structure in EVMS that is capable of being a block device. Disks, segments, regions, and feature objects are all storage objects.
Container	A collection of storage objects that are combined to form a new storage object similar to volume groups in LVM.
Volume (Logical Volume)	A mountable storage object.
EVMS Volume (EVMS Logical Volume)	A mountable storage object that contains EVMS metadata that gives the volume a user-specified name.
Compatibility Volume (Compatibility Logical Volume)	<p>A mountable storage object that does not contain EVMS metadata. Many plug-ins in EVMS provide support for the capabilities of other volume management schemes. Volumes that are designated as Compatibility are insured to be backward compatible to that particular scheme because they do not contain EVMS metadata.</p> <p>For example, this includes volumes that are mounted under the <code>/dev</code> directory by LVM, instead of the <code>/dev/evms</code> directory by EVMS.</p>

1.5 Mount Location of Storage Objects with EVMS

All devices that EVMS manages are mounted under `/dev/evms`. For example:

Table 1-3 Mount Location of EVMS Storage Objects

Storage Object	Mount Location
A disk segment of disk <code>sda</code> called <code>/dev/sda5</code>	<code>/dev/evms/sda5</code>
A software RAID device called <code>/dev/md1</code>	<code>/dev/evms/md/md1</code>
An LVM volume called <code>/dev/lvm_group/lvm_volume</code>	<code>/dev/evms/lvm/lvm_group/lvm_volume</code>

Using EVMS to Manage Devices

2

This section describes how to configure EVMS as the volume manager of your devices.

- [Section 2.1, “Configuring the System Device at Install to Use EVMS,” on page 15](#)
- [Section 2.2, “Configuring the System Device After the Install to Use EVMS,” on page 22](#)
- [Section 2.3, “Configuring LVM Devices After Install to Use EVMS,” on page 28](#)
- [Section 2.4, “Using EVMS with iSCSI Volumes,” on page 28](#)
- [Section 2.5, “Using the ELILO Loader Files \(IA-64\),” on page 29](#)
- [Section 2.6, “Starting EVMS,” on page 29](#)
- [Section 2.7, “Starting the EVMS UI,” on page 29](#)

2.1 Configuring the System Device at Install to Use EVMS

This section discusses how to configure the system device during the Linux install to use EVMS as the volume manager instead of the current default of Linux Volume Manager (LVM).

- [Section 2.1.1, “Before the Install,” on page 15](#)
- [Section 2.1.2, “During the Install,” on page 17](#)
- [Section 2.1.3, “After the Install,” on page 20](#)

2.1.1 Before the Install

- [“System Device” on page 15](#)
- [“Device Size Limits” on page 16](#)
- [“Data Loss Considerations for the System Device” on page 16](#)
- [“Storage Deployment Considerations for the System Device” on page 16](#)

System Device

For the purposes of this install documentation, a system device is any device that contains the Linux `/boot`, `swap`, or `root (/)` partitions for your Linux computer.

The install instructions assume the following:

- All three system partitions are on the same physical disk.

If you use different disks for any of the system partitions, make sure to modify the install instructions for your deployment scenario so that all of the system partitions are managed by EVMS.

- You must configure the boot partition within the BIOS-addressable space (such as 2 GB for x86 or 8 GB for x86-64) of the first disk recognized by the system.

If this restriction is not required for your hardware, you can modify the location of the `/boot` partition as desired.

- Your system uses the Grub or LILO boot loaders.

If you have an IA64 system, you must modify these install instructions to use the ELILO boot loader (`/boot/efi/elilo.conf`) instead.

WARNING: Whenever you manually alter the kernel or `initrd` on your system, make sure to run `/sbin/elilo` before shutting down the computer. If you leave out this step, your system might not be bootable.

Device Size Limits

The MS-DOS* partition table format does not recognize block devices larger than 2 TB. Make sure that the partitions you create for the `/boot` and `root (/)` are on devices up to 2 TB in size.

IMPORTANT: If you have a local disk, external disk arrays, or SAN devices that are larger than 2 TB, use a third-party disk partitioner to carve the devices into logical devices that are less than 2 TB each.

Because the `md` driver and `mdadm` work with the MS-DOS partition table format, each member device of a software RAID can be only up to 2 TB. The `md` RAID device itself can be larger than 2 TB and up to the maximum device size supported by the file system you plan to use. For information about file system limits for SUSE® Linux Enterprise Server 10, see “Large File System Support” in the *SUSE Linux Enterprise Server 10 Installation and Administration Guide*. (<http://www.novell.com/documentation/sles10>).

Data Loss Considerations for the System Device

This install requires that you delete the default *Partitioning* settings created by the install, and create new partitions to use EVMS instead. This destroys all data on the disk.

WARNING: To avoid data loss, it is best to use the EVMS install option only on a new device.

If you have data volumes on the system device, take one or more of the following precautionary measures:

- Move the data volumes from the system device to another device.
- If you cannot move the volumes, make a backup copy of the data, so you can restore the data volumes later from a backup copy.

Storage Deployment Considerations for the System Device

By default, the YaST install for SUSE Linux Enterprise Server uses the Linux Volume Manager to manage the system device. The install procedures in this section describe how to install SUSE Linux Enterprise Server with EVMS as the volume manager of the system device. The instructions assume the following:

- You want to use EVMS to manage the system device.
- Only the system device is to be configured during the install.
- Other devices on the system are not configured during the install, or are attached to the server post-install. These additional devices are configured only after the system is operating and performing as expected.

2.1.2 During the Install

WARNING: The following install destroys all data on the system device.

To install Linux with EVMS as the volume manager for your boot and system partitions, you must modify the Partitioning configuration in the *Installation Settings*.

- 1** Begin the install, according to the instructions provided in the “Deployment” section of the *SUSE Linux Enterprise 10 Administration Guide*.
- 2** When the install process reaches the *Installations Settings* screen, delete the recommended partitions and the partition table on the system disk so that the device can be marked to use EVMS as the volume manager instead of LVM.
 - 2a** In the list of *Installation Settings*, select *Partitioning*.
 - 2b** In the *Partitioning* menu, select *Create Custom Partition Setup*, then click *Next*.
 - 2c** Select *Custom Partition - for Experts*, then click *Next* to open the Expert Partitioner options.
 - 2d** Select *Expert > Delete Partition Table and Disk Label*, then click *Yes* twice to continue through the Warning advisories.

This deletes the recommended partitions and the partition table on the system disk.

- 3** Create a primary partition on the system disk to use as the boot partition.

During the install, the `/boot` partition must remain under LVM control so that the install completes successfully. You do not modify the volume manager of the `/boot` partition until after the install is complete.

- 3a** Click *Create*.
- 3b** From the list of devices, select the device you want to use for the boot partition, such as `/dev/sda`, then click *OK*.

If you have a single system disk, only one device is available and you are not prompted for the device.
- 3c** Select *Primary Partition*, then click *OK*.
- 3d** Select *Format*, then select the native Linux file system you want to use, such as Reiser.
- 3e** In *Size (End Value)* field, specify 200 MB or larger.

For example, to set the size at 300 MB, type `300M`.
- 3f** Set the mount point to `/boot`.
- 3g** Click *OK*.

The partition appears as a logical device in the devices list, such as `/dev/sda1`.

- 4** Create a second primary partition on the system disk where you later create your swap and system volumes:

- 4a** Click *Create*.
- 4b** From the list of devices, select the device you want to use for the second primary partition, such as `/dev/sda`, then click *OK*.

If you have a single system disk, only one device is available and you are not prompted for the device.
- 4c** Select *Primary Partition*, then click *OK*.

- 4d** Select *Do Not Format*, then select *Linux LVM (0x8E)* from the list of file system IDs.
- 4e** In *Size*, set the cylinder *End Value* to 5 GB or larger, depending on the combined partition size you need to contain your system and swap volumes. You are creating a primary partition that becomes the EVMS container for these two volumes that you create later in this procedure.

IMPORTANT: Do not make the system partition larger than necessary. The remaining space on the system disk can be used to create data volumes that are managed by EVMS.

To determine the size you need, consider the following recommendations:

- If you intend to create data volumes on the same physical disk, you must leave unpartitioned space available.
- Set aside 128 MB or larger for the swap partition.

Swap management is different for Linux kernel 2.4.10 and later. How much swap to add depends on the RAM size, the tasks that are planned for the system, and whether you want to make more virtual memory available than the RAM provides.

Some swap (at least 128 MB) is good in order to minimize the risk of losing data when active processes run out of RAM space. Swap is not required for systems with more than 1 GB of RAM. You must have at least 1 GB of virtual memory (RAM plus swap) during the install, but if the swap is more than 2 GB, you might not be able to install on some machines.

- 4f** Click *OK*.

The partition appears as a logical device in the devices list, such as `/dev/sda2`.

- 5** Modify the volume management type from LVM to EVMS for the second primary partition you created in [Step 4](#) as follows:

- 5a** At the bottom of the page, click *EVMS*.

- 5b** In the EVMS Configuration dialog box, click *Create Container*, then select the LVM partition created in [Step 4](#).

- 5c** In the Create EVMS Container dialog box, specify the container name (such as *system*), then click *Add Volume* to create the `lvm2/system` container, where *system* is the container name.

- 5d** Click *OK*.

The EVMS Configuration dialog box now displays the `lvm2/system` container you just created, its size, and free space.

- 6** In the EVMS Configuration dialog box, create the *swap* volume in the `lvm2/system` container as follows:

- 6a** From the *EVMS Container* drop-down menu, select `lvm2/system`, then click *Add*.

- 6b** In the Create Logical Volume dialog box, select *Format*, then select *Swap* from the *File System* drop-down menu.

- 6c** Specify *swap* as the volume name.

- 6d** Specify the size of the *swap* volume, such as 1GB.

The *swap* volume should be at least 128 MB or larger. For more information, see [Step 4e on page 18](#).

- 6e** Specify the mount point as *swap*.

6f Click *OK*.

The swap volume is now listed as a volume in the `lvm2/system` container.

- 7** In the EVMS Configuration dialog box, create the root (`/`) volume in the `lvm2/system` container as follows:

7a From the *EVMS Container* drop-down menu, select `lvm2/system`, then click *Add*.

7b In the Create Logical Volume dialog box, select *Format*, then select the file system to use from the *File System* drop-down menu, such as *Reiser*.

7c In the *Volume Name* field, specify a volume name, such as `sys_lx`.

7d Specify the *Size* of the system volume as the remaining space available in the `lvm2/system` partition by clicking *Max*.

7e Specify the mount point as `/` (root volume).

7f Click *OK*.

The root (`/`) volume is now listed as a volume in the `lvm2/system` container.

- 8** Click *Next* to return to the list of devices.

Below is an example of the physical and logical devices that might be configured on your system. Your setup depends on the number of devices in the server and the sizes you choose for your partitions.

Device	Size	F	Type	Mount	Start	End	Used By
<code>/dev/sda</code>	149.0 GB		6Y160p0		0	19456	
<code>/dev/sda1</code>	305.9 MB	F	Linux Native (Reiser)	<code>/boot</code>	0	38	
<code>/dev/sda2</code>	20.0 GB		Linux LVM		39	2649	EVMS lvm2/ system
<code>/dev/sdb</code>	111.8 GB		SP1203N		0	14595	
<code>/dev/evms/lvm2/system/sys_lx</code>	10.0 GB	F	EVMS	<code>/</code>	-	-	
<code>/dev/evms/lvm2/system/swap</code>	1.0 GB	F	EVMS	<code>swap</code>	-	-	

- 9** Click *Next* to return to the Installation Settings page.

You can dismiss the message warning that you should not mix EVMS and non-EVMS partitions on the same device.

- 10** Continue with the installation.

IMPORTANT: After the install is complete, make sure to perform the mandatory post-install configuration of the related system settings to ensure that the system device functions properly under EVMS. Otherwise, the system fails to boot properly.

For information, see [“After the Install” on page 20](#).

2.1.3 After the Install

After the install is complete, you must perform the following tasks to ensure that the system device functions properly under EVMS:

- “Edit the `/etc/fstab` File” on page 20
- “Disable the `boot.lvm` and `boot.md` Services” on page 21
- “Enable the `boot.evms` Service” on page 21
- “Edit the `/etc/init.d/boot.evms` Script” on page 21
- “Reboot the Server” on page 22
- “Verify the System Services” on page 22

Edit the `/etc/fstab` File

When you boot the system, the kernel reads the `/etc/fstab` file to identify which file systems should be mounted and then mounts them. This file contains a table of file system information about the root (`/`), `/boot`, and swap partitions plus other partitions and file systems you want to mount.

The `/boot` partition is separate from the EVMS container where you placed the root (`/`) and swap partitions and is not managed by EVMS at this time. However, in the following steps, you disable `boot.lvm` and `boot.md`, then enable `boot.evms`. In effect, this forces EVMS to scan all the partitions at boot time, including the `/boot` partition, and it activates `/boot` under the `/dev/evms` directory. Therefore, this makes `/boot` a partition that is discovered by EVMS at startup, and requires that the device be listed under `/dev/evms` in the `fstab` file so it can be found when booting with `boot.evms`.

After the install, you must edit the `/etc/fstab` file to modify the location of the `/boot` partition so it is under the `/dev/evms` directory.

IMPORTANT: When working in the `/etc/fstab` file, do not leave any stray characters or spaces in the file. This is a configuration file, and it is highly sensitive to such mistakes.

1 Open the `/etc/fstab` file in a text editor.

2 Locate the line that contains the `/boot` partition.

For example, if your `/boot` partition uses device `sda1` and the *Reiser* file system, look for a line similar to this:

```
/dev/sda1 /boot reiser defaults 1 1
```

3 In the *Device Name* column, modify the location of the `/boot` partition from `/dev` to `/dev/evms` so it can be managed by EVMS. Modify only the device name by adding `/evms` to the path:

```
/dev/evms/sda1 /boot reiser defaults 1 1
```

4 Save the file.

The changes do not take affect until the server is restarted. Do not reboot at this time.

5 Continue with “Disable the `boot.lvm` and `boot.md` Services” on page 21.

Disable the boot.lvm and boot.md Services

Disable the `boot.lvm` and `boot.md` services so they do not run at boot time (runlevel B). EVMS now handles the boot.

- 1 In YaST, click *System > System Services (Runlevel) > Expert Mode*.
- 2 Select *boot.lvm*.
- 3 Click *Set/Reset > Disable the Service*.
- 4 Select *boot.md*.
- 5 Click *Set/Reset > Disable the Service*.
- 6 Click *Finish*, then click *Yes*.

The changes do not take affect until the server is restarted. Do not reboot at this time.

- 7 Continue with [“Enable the boot.evms Service” on page 21](#).

Enable the boot.evms Service

The `boot.evms` service should be enabled automatically after the install, but you should verify that it is enabled.

- 1 In YaST, click *System > System Services (Runlevel) > Expert Mode*.
- 2 Select *boot.evms*.
- 3 Click *Set/Reset > Enable the Service*.

The *B runlevel* option is automatically selected.

- 4 Click *Finish*, then click *Yes*.

The changes do not take affect until the server is restarted. Do not reboot at this time.

- 5 Continue with [“Edit the /etc/init.d/boot.evms Script” on page 23](#).

Edit the /etc/init.d/boot.evms Script

- 1 Open the `/etc/init.d/boot.evms` script in a text editor.
- 2 Add the following lines to the Stop section:

```
mount -n -o remount,rw /
echo -en "\nDeleting devices nodes"
rm -rf /dev/evms
mount -n -o remount,ro /
```

For example, the Stop section looks like this after the edit:

```
stop)
    echo -n "Stopping EVMS"
    mount -n -o remount,rw /
    echo -en "\nDeleting devices nodes"
    rm -rf /dev/evms
    mount -n -o remount,ro /
    rc_status -v
;;
```

- 3 Save the file.
- 4 Continue with [“Reboot the Server” on page 22.](#)

Reboot the Server

- 1 Reboot the server to apply the post-install configuration settings.

Verify the System Services

After the post-install configuration is complete and you have rebooted the server, make sure the server is operating as expected.

2.2 Configuring the System Device After the Install to Use EVMS

If you have already installed Linux with a different volume manager for the system device (that is, the devices where you installed the `/boot`, `swap`, or `root (/)` partitions), you can optionally configure the device for EVMS after the install.

If you do not configure the device to use EVMS, you must manage the device and all of its volumes with its current volume manager (the default is LVM), and free space on the device cannot be used for volumes you want to create using EVMS. Beginning with the Linux 2.6 kernel, any given device cannot be managed by multiple volume managers. However, you can have different volume managers for different devices.

The following procedures assume that you installed Linux with three partitions on a single SCSI device named `sda`:

```
/dev/sda1 reiserfs /boot
/dev/sda2 swap      swap
/dev/sda3 reiserfs /
```

IMPORTANT: Make sure to modify the following procedures as necessary for your specific setup.

- [Section 2.2.1, “Disable the boot.lvm and boot.md Services,” on page 23](#)
- [Section 2.2.2, “Enable the boot.evms Service,” on page 23](#)
- [Section 2.2.3, “Edit the /etc/init.d/boot.evms Script,” on page 23](#)
- [Section 2.2.4, “Edit the /etc/fstab File,” on page 24](#)
- [Section 2.2.5, “Edit the Boot Loader File,” on page 25](#)
- [Section 2.2.6, “Force the RAM Disk to Recognize the Root Partition,” on page 26](#)
- [Section 2.2.7, “Reboot the Server,” on page 27](#)
- [Section 2.2.8, “Verify that EVMS Manages the Boot, Swap, and Root Partitions,” on page 27](#)

2.2.1 Disable the boot.lvm and boot.md Services

You need to disable `boot.lvm` (handles devices for Linux Volume Manager) and `boot.md` (handles multiple devices in software RAIDs) so they do not run at boot time. In the future, you want `boot.evms` to run at boot time instead.

1 In YaST, click *System > Runlevel Editor > Expert Mode*.

2 Select *boot.lvm*.

3 Click *Set/Reset > Disable the Service*.

4 Select *boot.md*.

5 Click *Set/Reset > Disable the Service*.

6 Click *Finish*, then click *Yes*.

The changes do not take affect until the server is restarted. Do not reboot at this time.

7 Continue with [Section 2.2.2, “Enable the boot.evms Service,” on page 23](#).

2.2.2 Enable the boot.evms Service

You need to enable the `boot.evms` service so that it boots devices when you restart the server.

1 In YaST, click *System > Runlevel Editor > Expert Mode*.

2 Select *boot.evms*.

3 Click *Set/Reset > Enable the Service*.

The *B runlevel* option is automatically selected.

4 Click *Finish*, then click *Yes*.

The changes do not take affect until the server is restarted. Do not reboot at this time.

5 Continue with [“Edit the /etc/init.d/boot.evms Script” on page 23](#).

2.2.3 Edit the /etc/init.d/boot.evms Script

1 Open the `/etc/init.d/boot.evms` script in a text editor.

2 Add the following lines to the Stop section:

```
mount -n -o remount,rw /
echo -en "\nDeleting devices nodes"
rm -rf /dev/evms
mount -n -o remount,ro /
```

For example, the Stop section looks like this after the edit:

```
stop)
    echo -n "Stopping EVMS"
    mount -n -o remount,rw /
    echo -en "\nDeleting devices nodes"
    rm -rf /dev/evms
    mount -n -o remount,ro /
    rc_status -v
;;
```

3 Save the file.

The changes do not take affect until the server is restarted. Do not reboot at this time.

4 Continue with [Section 2.2.4, “Edit the /etc/fstab File,” on page 24](#).

2.2.4 Edit the /etc/fstab File

When you boot the system, the kernel reads the `/etc/fstab` file to identify which file systems should be mounted and then mounts them. This file contains a table of file system information about the `/boot`, `swap`, and `root (/)` partitions plus other partitions and file systems you want to mount.

You must edit the `/etc/fstab` file to modify the mount location of these three partitions so they are mounted under the `/dev/evms` directory. For example, change `/dev/sda1` to `/dev/evms/sda1`.

Although the `/boot` partition is not managed by EVMS, the `boot.evms` script forces EVMS to scan all the partitions at boot time, including the `/boot` partition, and it activates `/boot` under the `/dev/evms` directory. Therefore, this makes `/boot` a partition that is discovered by EVMS at startup, and requires that the device’s path be listed under `/dev/evms` in the `fstab` file so it can be found when booting with `boot.evms`.

Make sure to replace `sda1`, `sda2`, and `sda3` with the device names you used for your partitions.

IMPORTANT: When working in the `/etc/fstab` file, do not leave any stray characters or spaces in the file. This is a configuration file, and it is highly sensitive to such mistakes.

1 Open the `/etc/fstab` file in a text editor.

2 Locate the line that contains the `/boot` partition.

For example, if your `/boot` partition uses device `sda1` and the *Reiser* file system, look for a line similar to this:

```
/dev/sda1 /boot reiser defaults 1 1
```

3 In the *Device Name* column, modify the mount location of the `/boot` partition from `/dev` to `/dev/evms` so it can be managed by EVMS. Modify only the device name by adding `/evms` to the path:

```
/dev/evms/sda1 /boot reiser defaults 1 1
```

4 Repeat [Step 2](#) and [Step 3](#) to edit the Device Name entry in the lines for the `swap` and `root (/)` partitions.

For example, change `/dev/sda2` to `/dev/evms/sda2`, and change `/dev/sda3` to `/dev/evms/sda3`.

- 5 Save the file.

The changes do not take affect until the server is restarted. Do not reboot at this time.

- 6 Continue with [Section 2.2.5, “Edit the Boot Loader File,” on page 25.](#)

2.2.5 Edit the Boot Loader File

When you boot the system, the kernel reads the boot loader file for information about your system. For Grub, this is the `/boot/grub/menu.lst` file. For LILO, this is the `/etc/lilo.conf` file.

You must edit the boot loader file to modify the mount location of partitions so they are mounted under the `/dev/evms` directory. For example, change `/dev/sda1` to `/dev/evms/sda1`. Make sure to replace the path for all lines that contain device paths in the files. You can modify the boot loader file by editing fields in YaST, or use a text editor to modify the file directly.

IMPORTANT: When working in the boot loader file, do not leave any stray characters or spaces in the file. This is a configuration file, and it is highly sensitive to such mistakes.

Using YaST

To modify the boot loader file in the YaST Control Center:

- 1 Log in as the `root` user or equivalent.
- 2 In Yast, select *System > Boot Loader*.
- 3 Modify the boot loader image so that the root file system is mounted as `/dev/evms/` instead of `/dev/`.
 - 3a Select the boot loader image file, then click *Edit*.
 - 3b Edit the device path in the *Root Device* field.

For example, change the *Root Device* value from

```
/dev/sda2
```

to

```
/dev/evms/sda2
```

Replace `sda2` with the actual device on your machine.
 - 3c Edit any device paths in the *Other Kernel Parameters* field.
 - 3d Click *OK* to save the changes and return to the Boot Loader page.
- 4 Modify the failsafe image so that the failsafe root file system is mounted as `/dev/evms/` instead of `/dev/`.
 - 4a Select the failsafe image file, then click *Edit*.
 - 4b Edit the device path in the *Root Device* field.
 - 4c Check the *Other Kernel Parameters* field and make changes if needed.
 - 4d Click *OK* to save the change and return to the Boot Loader page.
- 5 Click *Finish*.
- 6 Continue with [Section 2.2.6, “Force the RAM Disk to Recognize the Root Partition,” on page 26.](#)

Using a Text Editor

To edit the boot loader file in a text editor:

- 1 Log in as the `root` user or equivalent.
- 2 Open the boot loader file in a text editor.
For Grub, this is the `/boot/grub/menu.1st` file. For LILO, this is the `/etc/lilo.conf` file.
- 3 Locate the line that contains the `root=` parameter.
For example, if your root file system uses device `sda1`, look for a line similar to this:

```
kernel (sd0,0)/vmlinuz root=/dev/sda1 vga=0x31a splash=silent showopts
```
- 4 Modify the mount location from `/dev` to `/dev/evms` so it can be managed by EVMS.
For example, after the change, the line looks like this:

```
kernel (sd0,0)/vmlinuz root=/dev/evms/sda1 vga=0x31a splash=silent showopts
```
- 5 Repeat [Step 3](#) and [Step 4](#) to locate other lines in the file that need to be similarly modified.
- 6 Save the file.
The changes do not take affect until the server is restarted. Do not reboot at this time.
- 7 Continue with [Section 2.2.6, “Force the RAM Disk to Recognize the Root Partition,”](#) on [page 26](#).

2.2.6 Force the RAM Disk to Recognize the Root Partition

Although you modified the `/etc/fstab` file so that everything now mounts from the `/dev/evms` path, the RAM disk still tries to load the root (`/`) partition from the `/dev` path. You must force the RAM disk to update its information.

- 1 At a terminal prompt, enter the EVMS Ncurses command as the `root` user or equivalent:

```
evmsn
```
- 2 Review the output to verify that EVMS shows only the `/boot` and `swap` partitions as active in EVMS.
You should see the following devices mounted (with your own partition names, of course) for these two partitions:

```
/dev/evms/sda1  
/dev/evms/sda2
```
- 3 Edit `/sbin/mkinitrd` to change line 2513 from

```
root_evms=
```


to

```
root_evms=1
```


then save the file.
- 4 At a terminal prompt, enter

```
/sbin/mkinitrd
```

This forces `initrd` to get the updated `/dev/evms` path information for the root (`/`) partition.

- 5 Edit `/sbin/mkinitrd` to change line 2513 from

```
root_evms=1
```

to

```
root_evms=
```

then save the file.

The changes do not take affect until the server is restarted. You reboot in the next task.

- 6 Continue with [Section 2.2.7, “Reboot the Server,”](#) on page 27.

2.2.7 Reboot the Server

- 1 Reboot the server to apply the post-install configuration settings.

When your system reboots, the kernel loads the `init-ramdisk`, which runs the EVMS tools to activate your volumes and mount your root file system. Then your boot scripts run the EVMS tools once more to make sure your `/dev/evms/` directory correctly reflects the current state of your volumes. Finally, the remaining EVMS volumes are mounted as specified in your `/etc/fstab` file. Everything else on your system should start up as you would normally expect.

- 2 Continue with [Section 2.2.8, “Verify that EVMS Manages the Boot, Swap, and Root Partitions,”](#) on page 27.

2.2.8 Verify that EVMS Manages the Boot, Swap, and Root Partitions

- 1 At a terminal prompt, enter the EVMS Nurses command as the `root` user or equivalent:

```
evmsn
```

- 2 Review the output to verify that EVMS shows the `/boot`, `swap`, and `root (/)` partitions as active in EVMS.

You should see the following devices mounted (with your own partition names, of course) under the `/dev/evms` directory:

```
/dev/evms/sda1
```

```
/dev/evms/sda2
```

```
/dev/evms/sda3
```

2.3 Configuring LVM Devices After Install to Use EVMS

Use the following procedure to configure data devices (not system devices) to be managed by EVMS. If you need to configure an existing system device for EVMS, see [Section 2.2](#), “Configuring the System Device After the Install to Use EVMS,” on page 22.

- 1 In a terminal console, run the EVMSGUI by entering the following as the `root` user or equivalent:

```
evmsgui
```

- 2 In the *Volumes* panel, review the names that EVMS reports as compatibility volumes, find the devices that represent the devices you want to manage using EVMS, then write down the names for future reference.

For example, `/dev/sdb1`.

- 3 In a text editor, edit the `/etc/fstab` file to use the EVMS volume names.

For example, change the following entry for an LVM2 volume from this

```
/dev/sdb1 / reiserfs defaults 1 2
```

to this

```
/dev/evms/lvm2/sdb1 / reiserfs defaults 1 2
```

IMPORTANT: Make sure not to leave any stray characters or spaces in the line.

With these changes, each time your system boots, your file system is mounted using EVMS as the volume manager.

- 4 Update the boot scripts as follows:
 - The command `evms_activate` must be run from your boot scripts in order to activate your volumes so they can be mounted.
 - If you run software-RAID (`boot.md`) or LVM (`boot.lvm`) boot files in your boot scripts, and if you are moving all devices to EVMS, remove or disable those commands.
- 5 If you have not already done so, enable the `boot.evms` service.
For information, see [“Enable the boot.evms Service” on page 21](#).
- 6 Reboot your system.

2.4 Using EVMS with iSCSI Volumes

If your EVMS devices, RAID, and volumes use storage devices from an iSCSI SAN, make sure that your system starts iSCSI before EVMS so that the SAN and its disks are available to EVMS on system startup. iSCSI must be started and running before any disks/volumes on the iSCSI SAN can be accessed. If EVMS starts before iSCSI, EVMS cannot see or access the devices in the iSCSI SAN to mount the storage objects they contain, so the EVMS devices, RAID, and volumes might not be visible or accessible.

If EVMS starts before iSCSI on your system so that your EVMS devices, RAIDs, and volumes are not visible or accessible, you must correct the order in which iSCSI and EVMS are started. Enter the `chkconfig` command at the Linux server console of every server that is part of your iSCSI SAN.

- 1 At a terminal console prompt, enter either

```
chkconfig evms on
```

or

```
chkconfig boot.evms on
```

This ensures that EVMS and iSCSI start in the proper order each time your servers reboot.

2.5 Using the ELILO Loader Files (IA-64)

On a SUSE Linux Enterprise Server boot device EFI System Partition, the full paths to the loader and configuration files are:

```
/boot/efi/SuSE/elilo.efi
```

```
/boot/efi/SuSE/elilo.conf
```

When configuring partitioning during the install on IA64 systems, set the file system type for the `/boot` partition to `vfat`, then choose *Fstab Options* and set the *Arbitrary* option value to `umask=077` to ensure that the partition is accessible only to administrators.

WARNING: Whenever you manually alter the kernel or `initrd` on your system, make sure to run `/sbin/elilo` before shutting down the computer. If you leave out this step, your system might not be bootable.

2.6 Starting EVMS

If EVMS does not start during the system boot, you must activate it manually.

- 1 Open a terminal console, then log in as the `root` user or equivalent.
- 2 At the terminal console prompt, enter

```
evms_activate
```

2.7 Starting the EVMS UI

Use the following procedure to start the EVMS UI.

IMPORTANT: When you are done, make sure to exit the EVMS UI tool. When it is running, the EVMS UI tool locks the EVMS engine, potentially blocking other EVMS actions from taking place.

- 1 Open a terminal console, then log in as the `root` user or equivalent.
- 2 Enter one of the following commands to open the desired EVMS UI:

Command	Description
evmsgui	Starts the graphical interface for EVMS GUI. For information about features in this interface, see "EVMS GUI" (http://evms.sourceforge.net/user_guide/#GUI) in the <i>EVMS User Guide</i> at the EVMS project on SourceForge.net.
evmsn	Starts the text-mode interface for EVMS Ncurses. For information about features in this interface, see the "EVMS Ncurses Interface" (http://evms.sourceforge.net/user_guide/#NCURSES) in the <i>EVMS User Guide</i> at the EVMS project on SourceForge.net.
evms	Starts the EVMS command-line interpreter (CLI) interface. For information about command options, see "EVMS Command Line Interpreter" (http://evms.sourceforge.net/user_guide/#COMMANDLINE) in the <i>EVMS User Guide</i> at the EVMS project on SourceForge.net.

Loading EVMS Devices by UUIDs

3

This section discusses the optional use of UUIDs instead of device names to identify devices in the boot loader file and the `/etc/fstab` file.

- [Section 3.1, “Naming Devices with udev,” on page 31](#)
- [Section 3.2, “Understanding UUIDs,” on page 31](#)
- [Section 3.3, “Using UUIDs in the Boot Loader and `/etc/fstab` File \(x86\),” on page 32](#)
- [Section 3.4, “Using UUIDs in the Boot Loader and `/etc/fstab` File \(IA64\),” on page 33](#)

3.1 Naming Devices with udev

In the Linux 2.6 and later kernel, `udev` provides a userspace solution for the dynamic `/dev` directory, with persistent device naming. As part of the hotplug system, `udev` is executed if a device is added or removed from the system.

A list of rules is used to match against specific device attributes. The `udev` rules infrastructure (defined in the `/etc/udev/rules.d` directory) provides stable names for all disk devices, regardless of their order of recognition or the connection used for the device. The `udev` tools examine every appropriate block device that the kernel creates to apply naming rules based on certain buses, drive types, or file systems. For information about how to define your own rules for `udev`, see *Writing udev Rules* (http://reactivated.net/writing_udev_rules.html), a document by Daniel Drake, the `udev` developer.

Along with the dynamic kernel-provided device node name, `udev` maintains classes of persistent symbolic links pointing to the device in the `/dev/disk` directory, which is further categorized by the `by-id`, `by-label`, `by-path`, and `by-uuid` subdirectories.

For more information about `udev` (8), see its man page. Enter the following at a terminal console prompt:

```
man 8 udev
```

3.2 Understanding UUIDs

A UUID (Universally Unique Identifier) is a 128-bit number that is unique on both the local system and across other systems. It is a randomly generated with system hardware information and time stamps as part of its seed. UUIDs are commonly used to uniquely tag devices.

- [Section 3.2.1, “Using UUIDs to Assemble or Activate Devices,” on page 31](#)
- [Section 3.2.2, “Finding the UUID for a Device,” on page 32](#)

3.2.1 Using UUIDs to Assemble or Activate Devices

The UUID is always unique to the partition and does not depend on the order in which it appears or where it is mounted. With certain SAN devices attached to the server, the system partitions are renamed and moved to be the last device. For example, if root (`/`) is assigned to `/dev/sda1` during the install, it might be assigned to `/dev/sdg1` after the SAN is connected. One way to

avoid this problem is to use the UUID in the boot loader and `/etc/fstab` files for the boot device.

A UUID never changes, no matter where the device is mounted, so it can always be found at boot. In a boot loader file, you typically specify the location of the device (such as `/dev/sda1` or `/dev/evms/sda1`) to mount it at system boot. The boot loader can also mount devices by their UUIDs and admin-specified volume labels. However, if you use a label and file location, you cannot change the label name when the partition is mounted.

You can use the UUID as criterion for assembling and activating software RAID devices. When a RAID is created, the `md` driver generates a UUID for the device, and stores the value in the `md` superblock.

3.2.2 Finding the UUID for a Device

You can find the UUID for any block device in the `/dev/disk/by-uuid` directory. For example, a UUID looks like this:

```
e014e482-1c2d-4d09-84ec-61b3aefde77a
```

3.3 Using UUIDs in the Boot Loader and `/etc/fstab` File (x86)

After the install, you can optionally use the following procedure to configure the UUID for the system device in the boot loader and `/etc/fstab` files for your x86 system.

- 1 Install the SUSE® Linux Enterprise Server for x86 with no SAN devices connected.
- 2 After the install, boot the system.
- 3 Open a terminal console as the `root` user or equivalent.
- 4 Navigate to the `/dev/disk/by-name` directory to find the UUID for the device where you installed `/boot`, `/root`, and `swap`.

4a At the terminal console prompt, enter

```
cd /dev/disk/by-name
```

4b List all partitions by entering

```
ll
```

4c Find the UUID, such as

```
e014e482-1c2d-4d09-84ec-61b3aefde77a -> /dev/sda1
```

- 5 Edit `/boot/grub/menu.lst` file, using the Boot Loader option in YaST2 or using a text editor.

For example, change

```
kernel /boot/vmlinuz root=/dev/sda1
```

to

```
kernel /boot/vmlinuz root=/dev/disk/by-uuid/e014e482-1c2d-4d09-84ec-61b3aefde77a
```

IMPORTANT: Make a copy of the original boot entry, then modify the copy. If you make a mistake, you can boot the server without the SAN connected, and fix the error.

If you use the Boot Loader option in YaST, there is a defect where it adds some duplicate lines to the boot loader file when you change a value. Use an editor to remove the following duplicate lines:

```
color white/blue black/light-gray
default 0
timeout 8
gfxmenu (sd0,1)/boot/message
```

When you use YaST to change the way that the root (/) device is mounted (such as by UUID or by label), the boot loader configuration needs to be saved again to make the change effective for the boot loader.

- 6** As the `root` user or equivalent, do one of the following to place the UUID in the `/etc/fstab` file:

- Open YaST to *System > Partitioner*, select the device of interest, then modify *Fstab Options*.
- Edit the `/etc/fstab` file to modify the system device from the location to the UUID.

For example, if the root (/) volume has a device path of `/dev/sda1` and its UUID is `e014e482-1c2d-4d09-84ec-61b3aefde77a`, change line entry from

```
/dev/sda1 / reiserfs acl,user_xattr 1 1
```

to

```
UUID=e014e482-1c2d-4d09-84ec-61b3aefde77a / reiserfs
acl,user_xattr 1 1
```

IMPORTANT: Make sure to make a backup copy of the `fstab` file before you begin, and do not leave stray characters or spaces in the file.

3.4 Using UUIDs in the Boot Loader and `/etc/fstab` File (IA64)

After the install, use the following procedure to configure the UUID for the system device in the boot loader and `/etc/fstab` files for your IA64 system. IA64 uses the EFI BIOS. Its file system configuration file is `/boot/efi/SuSE/elilo.conf` instead of `/etc/fstab`.

- 1** Install the SUSE Linux Enterprise Server for IA64 with no SAN devices connected.
- 2** After the install, boot the system.
- 3** Open a terminal console as the `root` user or equivalent.
- 4** Navigate to the `/dev/disk/by-name` directory to find the UUID for the device where you installed `/boot`, `/root`, and `swap`.

- 4a** At the terminal console prompt, enter

```
cd /dev/disk/by-name
```

- 4b** List all partitions by entering

```
ll
```

- 4c** Find the UUID, such as

```
e014e482-1c2d-4d09-84ec-61b3aefde77a -> /dev/sda1
```

- 5 Edit `/boot/grub/menu.lst` file, using the Boot Loader option in YaST2 or using a text editor.

For example, change

```
root=/dev/sda1
```

to

```
root=/dev/disk/by-uuid/e014e482-1c2d-4d09-84ec-61b3aefde77a
```

IMPORTANT: Make a copy of the original boot entry, then modify the copy. If you make a mistake, you can boot the server without the SAN connected, and fix the error.

If you use the Boot Loader option in YaST, there is a defect where it adds some duplicate lines to the boot loader file when you change a value. Use an editor to remove the following duplicate lines:

```
timeout = 80
```

```
read-only
```

```
relocatable
```

```
default = sles10
```

```
prompt
```

- 6 Edit the `/etc/fstab` file to modify the system device from the location to the UUID.

For example, change

```
/dev/sda1    /                    reiserfs    acl,user_xattr      1 1
```

to

```
UUID=e014e482-1c2d-4d09-84ec-61b3aefde77a    /    reiserfs
acl,user_xattr                                1 1
```

IMPORTANT: Make sure to make a backup copy of the `/etc/fstab` file before you begin, and do not leave stray characters or spaces in the file.

Managing Devices

4

This section discusses how to initialize a disk by adding a segment management container to manage the partitions that you later add to the disk.

- [Section 4.1, “Understanding Disk Segmentation,” on page 35](#)
- [Section 4.2, “Initializing Disks,” on page 36](#)
- [Section 4.3, “Removing the Segment Manager from a Device,” on page 38](#)
- [Section 4.4, “Creating Disk Segments \(or Partitions\),” on page 38](#)
- [Section 4.5, “Configuring Mount Options for Devices,” on page 39](#)
- [Section 4.6, “What’s Next,” on page 41](#)

4.1 Understanding Disk Segmentation

In EVMS, you initialize a disk by assigning a segment manager to it. The segment manager creates metadata for the disk and exposes its free space so you can subdivide it into one or multiple segments (also called partitions).

- [Section 4.1.1, “Segment Managers,” on page 35](#)
- [Section 4.1.2, “Disk Segments,” on page 36](#)

4.1.1 Segment Managers

The most commonly used segment manager is the DOS Segment Manager. The following table describes the segment managers available in EVMS.

Table 4-1 EVMS Segment Managers

Segment Manager	Description
DOS	The standard MS-DOS disk partitioning scheme. It is the most commonly used partitioning scheme for Linux, NetWare®, Windows*, OS/2*, BSD, SolarisX86, and UnixWare*.
GPT (Globally Unique Identifier (GUID) Partitioning Table)	<p>A partitioning scheme used for IA-64 platforms, as defined in the Intel* <i>Extensible Firmware Interface (EFI) Specification</i>. It is not compatible with DOS, Windows, or OS/2 systems.</p> <p>The GUID is also known as Universally Unique Identifier (UUID). The GPT combines time and space descriptors to create this unique 128-bit tag for the disk and its segments.</p>
Cluster	A partitioning scheme used exclusively for System/390 mainframes.
S/390	A partitioning scheme for high-availability clusters. It provides a GUID for the disk, creates an EVMS container for the shared cluster devices, and specifies a node ID for the node that owns the device and the cluster ID.
BSD	A partitioning scheme for BSD UNIX.

Segment Manager	Description
MAC	A partitioning scheme for Mac-OS partitions.

4.1.2 Disk Segments

After you initialize the disk by adding a segment manager, you see metadata and free space segments on the disk. You can then create one or multiple data segments in a disk segment.

Table 4-2 *Disk Segment Types*

Segment Type	Description
Metadata	A set of contiguous sectors that contain information needed by the segment manager.
Free Space	A set of contiguous sectors that are unallocated or not in use. Free space can be used to create a segment.
Data	A set of contiguous sectors that has been allocated from a disk. The segment might be in use for a volume or a software RAID.

4.2 Initializing Disks

You must initialize new disks and disks that you want to reformat. After the disk is initialized, you can subdivide, or carve, the device into one or more disk segments for your file systems.

- [Section 4.2.1, “Before You Begin,” on page 36](#)
- [Section 4.2.2, “Guidelines,” on page 37](#)
- [Section 4.2.3, “Adding a Segment Manager,” on page 37](#)

4.2.1 Before You Begin

If you use large disks or disk arrays, use the vendor’s tools to carve them into the sizes that are usable for the management tools you plan to use. For example, the md driver recognizes disks only up to 2 TB in size, so the limit also applies to the md plug-in for EVMS. Software RAID devices you create with EVMS can be larger than 2 TB, of course, because the md driver plug-in manages the disks underneath that storage structure.

When you boot the server, EVMS scans and recognizes all devices it manages. If you add a new device to the server or create a device using `mkfs`, EVMS automatically mounts it on reboot under `/dev/evms` as a compatibility volume, such as `/dev/evms/sdb`.

IMPORTANT: If you cannot find a new disk, device, or volume, look under `/dev/evms` in a file browser, or look for compatibility volumes in the Volumes Manager in the EVMS GUI (`evmsgui`).

4.2.2 Guidelines

Consider the following guidelines when initializing a disk:

- EVMS might allow you to create segments without first adding a segment manager for the disk, but it is best to explicitly add a segment manager to avoid problems later. You must add a segment manager if you plan to use the devices for volumes that you want to share as cluster resources.
- When you initialize an existing disk that is already formatted, the process of adding a segment manager destroys all data on the disk. If you want to keep the data on the disk, make sure to back up the data before you begin this process.
- For existing disks on the system or disks that you move from another system, you must delete any existing volume management structures, and remove any segment managers. This removes the device's metadata and data, and destroys all data on the disk.

WARNING: Do not initialize the device that contains your current system disk or any device that contains the `/boot`, `swap`, or `root (/)` volumes.

- If a new disk does not show up in the list of *Available Objects*, look for it in the *Volumes* list to see if the disk shows up as a compatibility volume. For example, a new disk `sdb` would show up as `/dev/evms/sdb`. Delete it from the *Volumes* list to force the disk to show up in *Available Objects*, then create segments as desired.

4.2.3 Adding a Segment Manager

Use the following procedure to assign a segment manager to device for servers using x86, x64, and IA64 controllers. This option is not available for S390 platforms, so simply continue with configuring software RAIDs or file system partitions, as desired.

WARNING: Adding a segment manager initializes the disk, completely removing all the segments it contains. All the data stored on the device is lost.

- 1 If the disk has any existing volume management structures or an existing segment manager, remove them.
 - 1a Select *Actions > Delete > Volume* to view the *Volumes* list.
 - 1b Select any existing volume management structures on the device, then click *Delete*.
 - 1c Select *Actions > Remove > Segment Manager from Storage Object*.
 - 1d Select the type of Segment Manager in use, then click *Next*.
 - 1e Select the device, then click *Remove*.
- 2 If the disk is a new one that is listed as a compatibility volume in the *Volumes* list, delete it as a compatibility volume.
 - 2a Select *Actions > Delete > Volume* to view the *Volumes* list.
 - 2b Select the device, then click *Delete*.
- 3 Add the Segment Manager.
 - 3a In the list of *Availability Objects*, select the device, then click *Actions > Add > Segment Manager to Storage Object*.

- 3b** From the list, select one of the following types of segment manager, then click *Next*.
 - *DOS Segment Manager* (the most common choice)
 - *GPT Segment Manager* (for IA-64 platforms)
 - *Cluster Segment Manager* (available only if it is a viable option for the selected disk)
- 3c** Select the device from the list of *Plugin Acceptable Objects*, then click *Next*.
- 3d** If required, specify the disk type as Linux.
- 3e** Click *Add* to create the segment management container for the disk, then click OK to dismiss the confirmation message.

4.3 Removing the Segment Manager from a Device

- 1** If the disk has any existing volume management structures, remove them.
 - 1a** Select *Actions > Delete > Volume* to view the Volumes list.
 - 1b** Select any existing volume management structures on the device, then click *Delete*.
- 2** Select *Actions > Remove > Segment Manager from Storage Object*.
- 3** Select the type of Segment Manager in use, then click *Next*.
- 4** Select the device, then click *Remove*.

4.4 Creating Disk Segments (or Partitions)

- 1** In EVMS, select *Actions > Create > Segment* to see a list of segment managers.
- 2** From the list, select the segment manager for the device you want to manage, then click *Next*.
 - *DOS Segment Manager* (the most common choice)
 - *GPT Segment Manager* (for IA-64 platforms)
 - *Cluster Segment Manager* (available only if it is a viable option for the selected disk)

For information about these and other segment managers available, see “[Segment Managers](#)” on page 35.

- 3** Select the storage object that you want to segment, then click *Next*.
- 4** Complete the required configuration options for the segment, and modify default values as desired.
 - *Size (MB)*: Specify the amount of space (in MB) that you want to use. Use the arrows or type in a value. The interface corrects the value to the lower or upper size limit if you specify a size that is too small or that exceeds the amount of free space available.
 - *Offset (sectors)*: Specify the number of sectors to skip before beginning this partition if you want to leave free space in front of it.
 - *Partition Type*: From the drop-down list, select *Linux* (default), *Linux Swap*, *Linux LVM*, *NTFS*, *HPFS*, *FAT16*, or *Other Partition Type*.
 - *Parti Type ID*: This value changes automatically based on the *Partition Type* value, except for the *Other Partition Type* option. Then you must manually enter a value.
 - *Bootable*: Click *Yes* to make a primary partition active so that you can boot from it, or click *No* to make it unbootable. No is the only option if you are creating a logical partition.

- *Primary Partition*: Click *Yes* for a primary partition, or click *No* for a logical partition.

Required settings are denoted in the page by an asterisk (*). All required fields must be completed to make the *Create* button active.

- 5 Click *Create* to create the segment.
- 6 Verify that the new segment appears in the Segment list.

4.5 Configuring Mount Options for Devices

The following table describes the *Fstab Options* that are configurable in YaST. The values are written to the `/etc/fstab` file and are applied upon reboot.

Table 4-3 *Fstab Options in YaST*

Fstab Option	Description
Mount by	<ul style="list-style-type: none"> • Device name (K) (default, such as <code>/dev/sda2</code>) • Volume label (L) Make sure to also specify a value in <i>Volume Label</i>. • UUID (U) For information about why you might want to discover partitions and devices by UUID, see Section 3.2.1, “Using UUIDs to Assemble or Activate Devices,” on page 31. • Device ID (I) • Device Path (P)
Volume label	A useful name to help you easily identify the volume on the server. By default, this field is empty.
Mount read-only	<p>Select the check box to enable this option. It is deselected (disabled) by default.</p> <p>If this option is enabled, files and directories cannot be modified or saved on the volume.</p>
No access time	<p>Select the check box to enable this option. It is deselected (disabled) by default.</p> <p>By default, the Linux <code>open(2)</code> command updates the access time whenever a file is opened. The <i>No Access Time</i> option disables the updating of access time, so that reading a file does not update its access time. Enabling the <i>No Access Time</i> option allows you to back up a volume without modifying the access times of its files.</p>
Mountable by user	<p>Select the check box to enable this option. It is deselected (disabled) by default.</p> <p>If this option is enabled, the volume can be mounted by any user; <code>root</code> privileges are not required.</p>
Do Not Mount at System Start-up	<p>Select the check box to enable this option. It is deselected (disabled) by default.</p> <p>The system volumes such as <code>/boot</code>, <code>swap</code>, and <code>root (/)</code> should all be mounted at system start. For other volumes, enable this option for a volume if you want to mount it manually later using the <code>mount</code> command at a terminal console prompt.</p>

Fstab Option	Description
Data journaling mode	<p>For journaling file systems, select the preferred journaling mode: ordered (default), journal, or writeback.</p> <ul style="list-style-type: none"> • Ordered: Writes data to the file system, then enters the metadata in the journal. • Journal: Writes data twice; once to the journal, then to the file system. • Writeback: Writes data to the file system and writes metadata in the journal, but the writes are performed in any order.
Access Control Lists (ACL)	Select this option to enable access control lists on the file system. It is enabled by default.
Extended user attributes	Select this option to enable extended user attributes on the file system. It is enabled by default.
Arbitrary option value	Specify any mount option that is legal for the Mount Options column for a device entry in the <code>/etc/fstab</code> file. Use a comma with no spaces to separate multiple options.

You can modify these values for each entry by editing the `/etc/fstab` file, or use the following procedure to modify the mount options for a volume in the `/etc/fstab` file from YaST.

- 1 Open YaST, then click *System > Partitioning*.
- 2 Select the device you want to modify, then click *Fstab Options*.

Mount in /etc/fstab By: Normally, a file system to mount is identified in `/etc/fstab` by the device name. This identification can be changed so the file system to mount is found by searching for a UUID or a volume label. Not all file systems can be mounted by UUID or a volume label. If an option is disabled, it is not possible.

Volume label: The name entered in this field is used as the volume label. This normally only makes sense when you activate the option for mounting by volume label. In the volume label is not possible to use character `/` and space.

Mount Read-Only: No writable access to the file system is possible. Default is false.

No access time: Access times are not updated when a file is read. Default is false.

Mountable by User: The file system may be mounted by an ordinary user. Default is false.

Not Mounted at System Start-up: The file system is not automatically mounted when the system starts. An entry in `/etc/fstab` is created and the file system is mounted with the appropriate options when the command `mount <mount point>` (`<mount point>` is the directory to which the

Fstab options:

Mount in /etc/fstab by

☐ Device name ☐ Device ID

☐ Volume label ☐ Device Path

☒ UUID

Volume Label

Root

☐ Mount read-only

☐ No access time

☐ Mountable by user

☐ Do Not Mount at System Start-up

Data Journaling Mode

ordered

☒ Access Control Lists (ACL)

☒ Extended User Attributes

Arbitrary option value

OK Cancel

3 Modify the settings as desired, then click *OK* to accept your changes.

4.6 What's Next

If multiple paths exist between your host bus adapters (HBAs) and the storage devices, configure multipathing for the devices before creating software RAIDs or file system volumes on the devices. For information, see [Chapter 5, “Managing Multipathing for Devices and Software RAIDs,” on page 43](#).

If you want to configure software RAIDs, do it before you create file systems on the devices. For information, see [Chapter 6, “Managing Software RAIDs,” on page 55](#).

Managing Multipathing for Devices and Software RAIDs

5

This section discusses how to configure multiple paths between the servers and storage devices for automatic failover and optional load balancing.

- [Section 5.1, “Understanding Multipathing,” on page 43](#)
- [Section 5.2, “Before You Begin,” on page 46](#)
- [Section 5.3, “Adding multipathd to the Boot Sequence,” on page 48](#)
- [Section 5.4, “Starting Multipath I/O Services,” on page 48](#)
- [Section 5.5, “Configuring Multipath Settings in the /etc/multipath.conf File,” on page 48](#)
- [Section 5.6, “Resolving Stalled I/O,” on page 49](#)
- [Section 5.7, “Configuring Time-Out Settings for the HBA,” on page 49](#)
- [Section 5.8, “Configuring Multipath I/O for Devices,” on page 49](#)
- [Section 5.9, “Configuring Multipathing for an Existing Software RAID,” on page 51](#)
- [Section 5.10, “Adding Storage Dynamically for Use with Multipath I/O,” on page 53](#)
- [Section 5.11, “What’s Next,” on page 54](#)

5.1 Understanding Multipathing

- [Section 5.1.1, “What Is Multipathing?,” on page 43](#)
- [Section 5.1.2, “Benefits of Multipathing,” on page 44](#)
- [Section 5.1.3, “Guidelines for Multipathing,” on page 44](#)
- [Section 5.1.4, “Device Mapper Multipath I/O Module,” on page 45](#)
- [Section 5.1.5, “Multipath Tools,” on page 46](#)

5.1.1 What Is Multipathing?

Multipathing is the ability of a server to communicate with the same physical or logical device across multiple physical paths between host bus adapters in the server and the storage controllers for the device, typically in Fibre Channel (FC) or iSCSI SAN environments. You can also achieve multiple connections with direct attached storage for SCSI connections. Multipathing provides connection fault tolerance and can optionally provide load balancing across the available connections.

EVMS manages the multiple I/O paths for a device as a multipath I/O region. It requires the Device Mapper plug-in for EVMS. Each region consists of up to 8 multiple virtual paths, where a virtual path consists of the multiple hardware components that make up the connection.

Multipathing is managed at the device level. It is not possible to manage connections to Linux partitions contained in the device. You can combine multiple devices to create a software RAID or volumes as logical devices, but the physical device itself is managed at a lower level.

IMPORTANT: If you plan to use multipathing with software RAIDs, you should first configure multipathing for the devices to create the `mdN` storage object (where *N* is the assigned device number) for each device, then configure the RAID with them.

5.1.2 Benefits of Multipathing

Multipathing can help provide fault tolerance for the connection between the server and its storage devices, typically in a SAN configuration. When multipathing is configured and running, it automatically isolates and identifies device connection failures, and reroutes I/O around the failing components. A previously failed path is automatically reinstated when it becomes healthy again.

Typical connection problems involve faulty adapters, cables, or controllers. When you configure multipath I/O for a device, the multipath driver monitors the active connection between devices. When it detects I/O errors, the multipath driver fails over to a designated secondary path. When the primary path recovers, control is automatically returned to the primary connection.

5.1.3 Guidelines for Multipathing

Multipathing is available only under the following conditions:

- Multiple physical paths exist between host bus adapters in the server and host bus controllers for the storage device.
- Multipathing is not supported for the root (`/`) or `/boot` partitions because boot loaders do not recognize multiple paths and cannot handle them.
- For software RAID devices, multipathing runs underneath the software RAID, so multipathing should be configured prior to creating the software RAID devices. If you have already created a software RAID, additional steps are required to configure the multiple paths to its member devices without destroying the data it contains. For information, see [Section 5.9, “Configuring Multipathing for an Existing Software RAID,” on page 51](#).
- The host bus adapters and controllers must be supported by Device Mapper and its multipath tools.

Currently, Qlogic*, Emulex*, and LSI are supported. Refer to the *SUSE Linux Enterprise Server Administration Guide* for the latest list of supported adapters.

- Device partitioning is completed prior to configuring multipathing.
- The storage subsystem you use on the multipathed device must support multipathing. Most storage subsystems should work; however, they might require an appropriate entry in the `DEVICES` variable in the `/etc/multipath.conf` file.

For a list of supported storage subsystems that allows multiple paths to be detected automatically, see “10.1 Supported Hardware” in the *SUSE Linux Enterprise Server Administration Guide*.

- When configuring devices for multipathing, use the device names in the `/dev/disk/by-id` directory instead of the default device names (such as `/dev/sd*`), because the `/dev/disk/by-id` names persist over reboots.

5.1.4 Device Mapper Multipath I/O Module

The Device Mapper Multipath I/O (DM-MPIO) module provides the multipathing capability for Linux. It supports up to 8 paths per physical storage device. Multipath protects against failures in the paths to the device, and not failures in the device itself. If one of the paths is lost (for example, a network adapter breaks or a fiber-optic cable is removed), I/O will be redirected to the remaining paths. If an active path fails, the DM continues to balance traffic across the healthy paths. If all active paths fail, inactive secondary paths must be waked up, so failover occurs with a delay of approximately 30 seconds.

Table 5-1 *Multipath I/O Features*

Features	Description
Active/passive failover	If the storage array has multiple controllers, and only one controller is active at a time, then only the paths from the host to the active storage controller are active. Connections to the second and subsequent controllers are passive.
Active/active failover	If the storage array has multiple controllers that are concurrently active, all connections from the host to the controllers are active.
Load balancing	The Device Mapper driver automatically load balances traffic across all active paths.
Controller failover	When the active controller fails over to the passive, or standby, controller, the Device Mapper driver automatically activates the paths between the host and the standby, making them the primary paths. When the failed primary controller is reactivated as primary, the Device Mapper driver automatically activates the previously-downed paths, too.
Boot/Root device support	Multipathing is not supported for the <code>/boot</code> and <code>/</code> devices.

Multiple paths to a device are automatically detected under the device driver's path. For example, for the `cciss` driver, dual paths might be named as follows:

```
/dev/cciss/c1d0p2  
/dev/cciss/c2d0p2
```

The `mdadm` device names take the form `/dev/mdN`, where *N* is an autogenerated number for the device, beginning with 0 and issued sequentially as the devices are created, such as `/dev/md0`, `/dev/md1`, and so on.

If multiple paths are not automatically detected, you can configure them manually in the `/etc/multipath.conf` file.

5.1.5 Multipath Tools

The `multipath-tools` user-space package takes care of automatic path discovery and grouping. It automatically tests the path periodically, so that a previously failed path is automatically reinstated when it becomes healthy again. This minimizes the need for administrator attention in a production environment.

The tools are described in the following table:

Table 5-2 Tools Available in the `multipath-tools` Package

Tool	Description
<code>multipath</code>	Scans the system for multipathed devices and assembles them.
<code>multipathd</code>	Waits for maps events then executes <code>multipath</code> .
<code>devmap-name</code>	Provides a meaningful device name to <code>udev</code> for device maps (devmaps).
<code>kpartx</code>	Maps linear devmaps to partitions on the multipathed device, which makes it possible to create multipath monitoring for partitions on the device.

For a list of files included in this package, see the [multipath-tools Package Description \(http://www.novell.com/products/linuxpackages/suselinux/multipath-tools.html\)](http://www.novell.com/products/linuxpackages/suselinux/multipath-tools.html).

5.2 Before You Begin

- [Section 5.2.1, “Preparing the SAN Devices,” on page 46](#)
- [Section 5.2.2, “Configuring `mdadm.conf` and `lvm.conf` for Multipathing,” on page 47](#)
- [Section 5.2.3, “Configuring Device Drivers in `initrd`,” on page 47](#)

5.2.1 Preparing the SAN Devices

Before configuring SAN devices for use on your server, do the following tasks to prepare the SAN devices, as necessary:

- Configure and zone the SAN with the vendor’s tools.
- Configure permissions for host LUNs on the storage arrays with the vendor’s tools.
- Install the Linux HBA driver module. Upon module installation, the host bus is scanned and any devices present in the SAN that are permissioned for the host are discovered and presented to the host for further configuration.

NOTE: Ensure that the HBA driver you are using does not have native multipathing enabled.

See the vendor’s specific instructions for more details.

- After the driver module is loaded, discover the device nodes assigned to specific array LUNs or partitions.

5.2.2 Configuring mdadm.conf and lvm.conf for Multipathing

The default `mdadm.conf` and `lvm.conf` files do not work properly with multipathed devices. By default, both `md` and `LVM2` scan only the physical devices, and they ignore any symbolic links to devices and Device Mapper multipath I/O devices. When managing DM-MPIO devices, you want the opposite behavior, so they ignore all physical devices, and scan only the devices listed in the `/dev/disk/by-id` directory.

To avoid problems, do the following:

- 1 Open the `/etc/mdadm.conf` file in a text editor, then modify the `Devices` variable to scan for devices in the `/dev/disk/by-id` directory, as follows:

```
DEVICES /dev/disk/by-id/*
```

5.2.3 Configuring Device Drivers in initrd

Make sure that the device drivers for the storage subsystem controllers load automatically by adding them to the `INITRD_MODULES` variable in the `/etc/sysconfig/kernel` file, then re-create the `initrd`.

- 1 Open the `/etc/sysconfig/kernel` file in a text editor, then add the controller's driver module information in the `INITRD_MODULES` variable.

For example, if you have a RAID controller that uses the `cciss` driver and a QLogic controller that uses the `qla2xxx` driver, the `INITRD_MODULES` entry would be as follows:

```
INITRD_MODULES="cciss qla2xxx"
```

- 2 Re-create the `INITRD` by entering the following as the `root` user or equivalent at a terminal console prompt:

```
mkinitrd
```

- 3 If you are using LILO as your boot loader, you must reinstall LILO by entering the following as the `root` user or equivalent at a terminal console prompt:

```
/sbin/lilo
```

- 4 In a text editor, set the following value in the `/etc/sysconfig/hotplug` file:

```
HOTPLUG_USE_SUBFS=no
```

- 5 In a text editor, add the controller driver modules to the `INITRD_MODULES` variable within the `/etc/sysconfig/kernel` file.

This configures `initrd` to automatically load the device drivers at boot time for the controllers attached to the multipathed devices. The boot scripts detect multipathed devices only if the modules for the respective controllers are loaded at boot time.

For example, if you use QLogic controllers, add `qla2xxx` to the list of drivers as follows:

```
INITRD_MODULES="cciss reiserfs qla2xxx"
```

- 6 Re-create the `INITRD` on your system to reflect this change by entering the following at a terminal console prompt:

```
mkinitrd
```

If you are using GRUB as a boot manager, you do not need to make any further changes. The driver is loaded within the `INITRD` upon the next reboot.

If you are using LILO as a boot manager, you must run it once to update the boot record.

5.3 Adding multipathd to the Boot Sequence

If you are using multipath IO services, add `multipathd` to the boot sequence, using one of the following methods:

- [Section 5.3.1, “YaST,” on page 48](#)
- [Section 5.3.2, “Command Line,” on page 48](#)

5.3.1 YaST

- 1 In YaST, click *System > System Services (Runlevel) > Simple Mode*.
- 2 Select *multipathd*, then click *Enable*.
- 3 Click *OK* to acknowledge the service startup message.
- 4 Click *Finish*, then click *Yes*.

The changes do not take affect until the server is restarted.

5.3.2 Command Line

- 1 Open a terminal console, then log in as the `root` user or equivalent.
- 2 Enter

```
insserv boot.multipath multipathd
```

5.4 Starting Multipath I/O Services

If the `boot.multipath` service is not started on system boot, do the following:

- 1 Open a terminal console, then log in as the `root` user or equivalent.
- 2 Enter

```
/etc/init.d/boot.multipath start  
/etc/init.d/multipathd start
```

5.5 Configuring Multipath Settings in the `/etc/multipath.conf` File

- 1 Open the `/etc/multipath.conf` file in a text editor.
- 2 Configure multipathing to queue I/O if all paths fail concurrently. Uncomment the `defaults` section and its ending bracket, then add the `defaults_features` setting, as follows:

```
defaults {  
    default_features "1 queue_if_no_path"  
}
```
- 3 After you start multipath I/O services, the paths are listed under `/dev/disk/by-name`. The default name is the WWN (World Wide Name) of the logical unit. You can override this behavior with `/var/lib/multipath/bindings` to use user-friendly names instead.

Uncomment the `user_friendly_names` variable and change its setting from No to Yes.

```
user_friendly_names yes
```

5.6 Resolving Stalled I/O

If all paths fail concurrently and I/O is queued and stalled, do one of the following:

- Enter the following command at a terminal console prompt:

```
dmsetup message <NAME> 0 fail_if_no_path
```

Replace `<NAME>` with the correct map name. This causes all queued I/O to fail and propagates the error to the calling application.

Reactivate queueing by entering the following command at a terminal console prompt:

```
dmsetup message <NAME> 0 queue_if_no_path
```

- Reboot the system.

5.7 Configuring Time-Out Settings for the HBA

When using multipath I/O, you want any host bus adapter (HBA) or cable failures to be reported faster than when single paths are used, so that multipathing can fail over to the healthy path. Configure time-out settings for your HBA to opt for failover sooner by modifying its options in the `/etc/modprobe.conf.local` file.

For example, for the QLogic 2xxx family of host bus adapters, the following setting are recommended:

```
options qla2xxx qlport_down_retry=1 ql2xfailover=0 ql2xretrycount=5
```

5.8 Configuring Multipath I/O for Devices

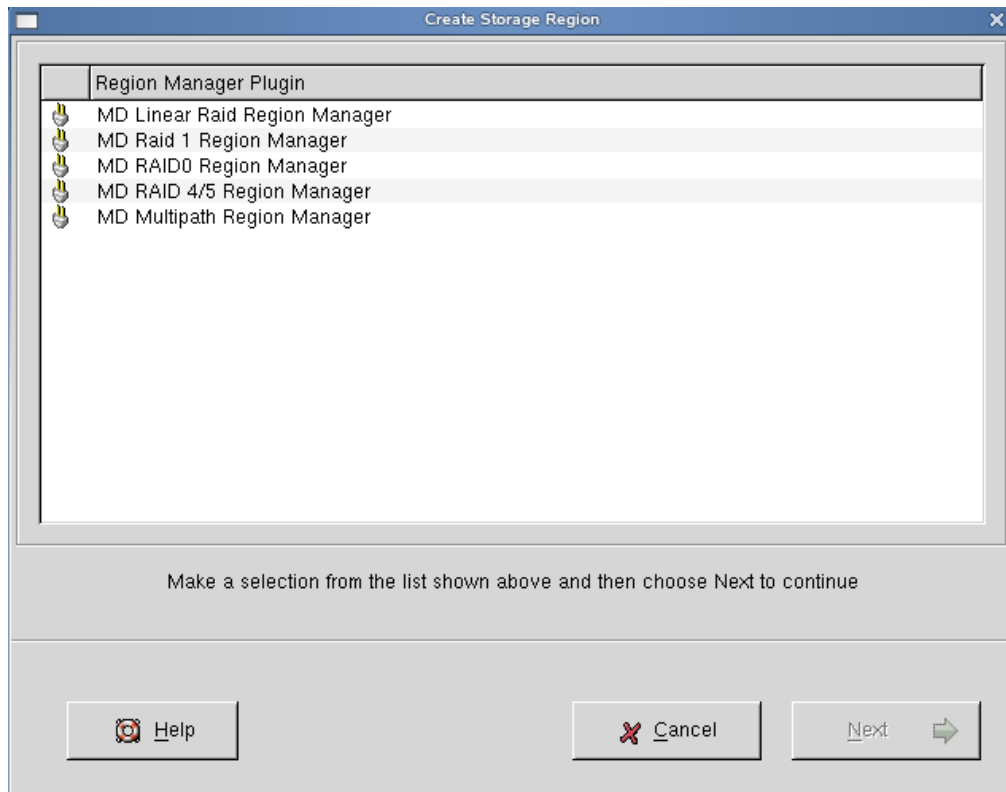
You should configure multipathing for storage devices before using them in a software RAID.

NOTE: If you have previously configured a software RAID device, see [Section 5.9, “Configuring Multipathing for an Existing Software RAID,”](#) on page 51.

- 1 Open a terminal console, then log in as the `root` user or equivalent.
- 2 Start the EVMS GUI tool by entering

```
evmsgui
```

- 3 Select *Action > Create > Region* to open the Create Storage Region dialog box.



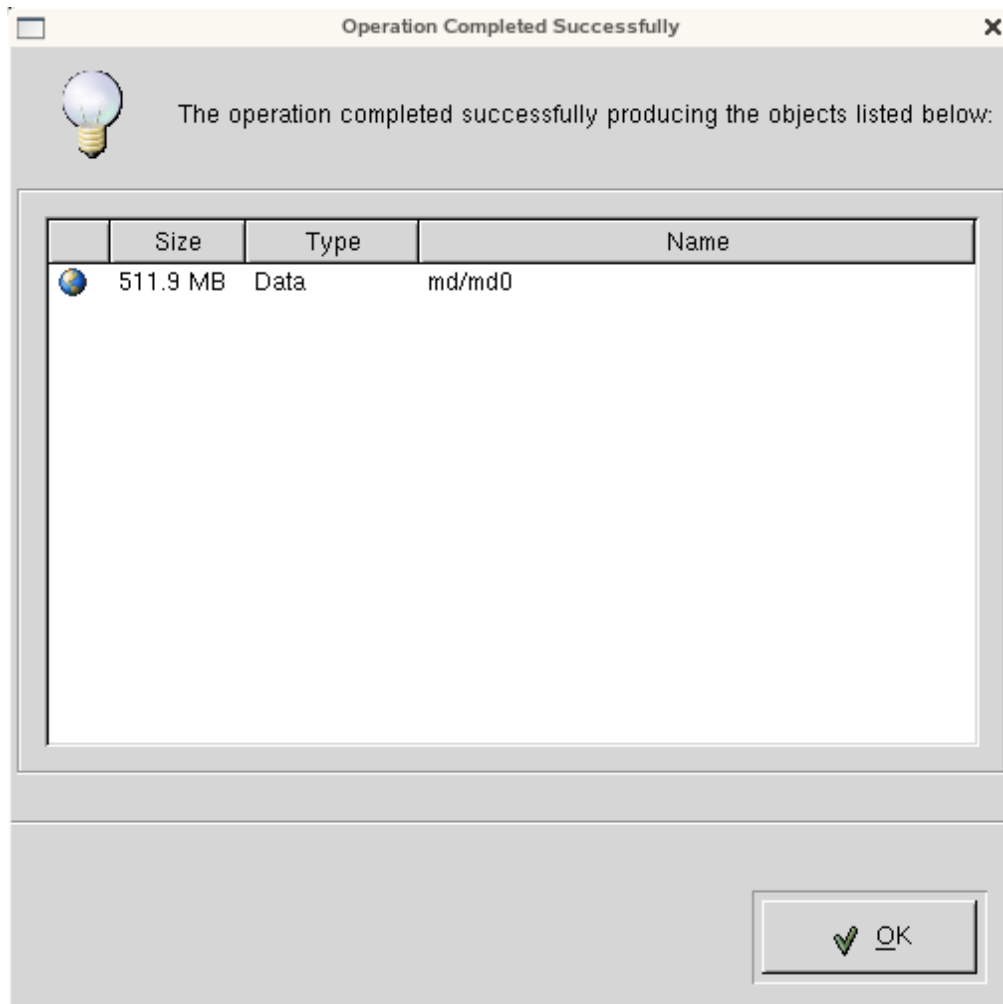
- 4 Select MD Multipath Region Manager, then click *Next*.
- 5 From the Storage Objects listed, select the ones that represent the same device.

IMPORTANT: The order of the objects for multipathing is implied by their order in the list.

- 6 Click *Create* to create the multipath device under the `/dev/evms/md` directory.

The device is given a name such as `md0`, so its EVMS mount location is `/dev/evms/md/md0`.

- 7 When the object is created, a message notifies you of its relative path name. Click *OK* to dismiss the message.



Now you can manage the device under its multipath identity of `/dev/evms/md/md0`, just as you would any for a device without multiple paths.

5.9 Configuring Multipathing for an Existing Software RAID

Ideally, you should configure multipathing for devices before you use them as components of a software RAID device. You can use the procedure in this section to get multipathing running for a previously existing software RAID. For example:

- If you create a new software RAID as part of the *Partitioning* settings during a new install or upgrade.
- If you did not configure the devices for multipathing before using them in the software RAID as a member device or spare.
- If you grow your system by adding new HBA adapters to the server or expanding the storage subsystem in your SAN.

If you add multipathing after creating any software RAID devices, the Multipath service might be starting after the MD service on reboot, which makes multipathing appear not to be available for RAIDs.

NOTE: The following instructions assume the software RAID device is `/dev/md0`, which is its device name as recognized by the kernel. Make sure to modify the instructions for the device name of your software RAID.

- 1** Open a terminal console, then log in as the `root` user or equivalent.

Except where otherwise directed, use this console to enter the commands in the following steps.
- 2** Start multipath services and enable them to start at reboot by entering the following commands:

```
chkconfig multipathd on
chkconfig boot.multipath on
```
- 3** Create and configure the `/etc/multipath.conf` file using the example in the `/usr/share/doc/packages/multipath-tools/multipath.conf.synthetic` file.
 - 3a** Enter the following command (all on one line, of course):

```
cp /usr/share/doc/packages/multipath-tools/
multipath.conf.synthetic /etc/multipath.conf
```
 - 3b** Open the `/etc/multipath.conf` file in a text editor.
 - 3c** Uncomment the `Defaults` directive and its ending bracket.
 - 3d** Uncomment the `user_friendly_names` option, then change its value from `No` to `Yes`.
 - 3e** Save your changes, then close the file.
- 4** Modify the `/etc/mdadm.conf` file to scan for devices in the `/dev/disk/by-id` directory.
 - 4a** Open the `/etc/mdadm.conf` file in a text editor.
 - 4b** Modify the `Devices` variable as follows:

```
DEVICES /dev/disk/by-id/*
```
 - 4c** Save your changes, then close the file.
- 5** If any software RAID devices are currently mounted or running, enter the following commands for each device to unmount the device and stop it.

```
umount /dev/md0
mdadm --misc --stop /dev/md0
```
- 6** Stop the `boot.md` service by entering

```
/etc/init.d/boot.md stop
```
- 7** Start the `boot.multipath` and `multipathd` services by entering the following commands:

```
/etc/init.d/boot.multipath start
/etc/init.s/multipathd start
```

- 8 After the multipathing services are started, verify that the software RAID's member devices are listed in the `/dev/disk/by-id` directory. Do one of the following:

- **Devices Are Listed:** The device names should now have symbolic links to their Device Mapper device names, such as `dev/dm-1`.
- **Devices Are Not Listed:** Force the `multipath` service to recognize them by flushing and rediscovering the devices.

To do this, enter the following commands:

```
multipath -F
multipath -v0
```

The devices should now be listed in `/dev/disk/by-id`, and have symbolic links to their Device Mapper device names. For example:

```
lrwxrwxrwx 1 root root 10 Jun 15 09:36 scsi-mpath1 -> ../../dm-1
```

- 9 Restart the `boot.md` service and the RAID device by entering

```
/etc/init.d/boot.md start
```

- 10 Check the status of the software RAID by entering

```
mdadm --detail /dev/md0
```

The RAID's member devices should match their Device Mapper device names that are listed as the symbolic links of devices in the `/dev/disk/by-id` directory.

- 11 Make a new `initrd` to ensure that the Device Mapper `multipath` services are loaded before the md RAID services on reboot. Enter

```
mkinitrd -f mpath
```

- 12 Reboot the server to apply these post-install configuration settings.

- 13 Verify that the software RAID array comes up properly on top of the multipathed devices by checking the RAID status. Enter

```
mdadm --detail /dev/md0
```

For example:

Number	Major	Minor	RaidDevice	State	
0	253	0	0	active sync	/dev/dm-0
1	253	1	1	active sync	/dev/dm-1
2	253	2	2	active sync	/dev/dm-2

5.10 Adding Storage Dynamically for Use with Multipath I/O

If your system has already been configured for multipathing and you later need to add more storage to the SAN, use the following procedure to rescan the devices and make them available to multipathing without rebooting the system.

- 1 On the storage subsystem, use the vendor's tools to allocate the devices and update its access control settings to allow the Linux system access to the new storage. Refer to the vendor's documentation for details.
- 2 On the Linux system, use the HBA driver commands to rescan the SAN to discover the new devices. The exact commands depend on the driver.

For example, for a QLogic 2300 HBA, the command is

```
echo scsi-qlascan >/proc/scsi/qla2xxx/<host number>
```

At this point, the newly added device is not known to the higher layers of the Linux kernel's SCSI subsystem and is not yet usable.

- 3 Rescan all targets for a host to make its new device known to the middle layer of the Linux kernel's SCSI subsystem. At a terminal console prompt, enter

```
echo "- - -" >/sys/class/scsi_host/host<hostnumber>/scan
```

- 4 Run the Multipath tool to recognize the devices for DM-MPIO configuration. At a terminal console prompt, enter

```
multipath
```

You can now configure the devices for multipathing.

5.11 What's Next

If you want to use software RAIDs, create and configure them before you create file systems on the devices. For information, see [Chapter 6, “Managing Software RAIDs,” on page 55](#).

Managing Software RAIDs

6

This section discusses how to manage devices and device connections with the Enterprise Volume Management System (EVMS).

- [Section 6.1, “Understanding Software RAIDs on Linux,” on page 55](#)
- [Section 6.2, “Creating and Configuring a Software RAID,” on page 61](#)
- [Section 6.3, “Expanding a RAID,” on page 65](#)
- [Section 6.4, “Adding or Removing a Spare Disk,” on page 66](#)
- [Section 6.5, “Removing a Failed Disk,” on page 67](#)
- [Section 6.6, “Replacing a Failed Disk,” on page 68](#)
- [Section 6.7, “Monitoring Status for a RAID,” on page 68](#)
- [Section 6.8, “Deleting a Software RAID and Its Data,” on page 72](#)

6.1 Understanding Software RAIDs on Linux

- [Section 6.1.1, “What Is a Software RAID?,” on page 55](#)
- [Section 6.1.2, “Overview of RAID Levels,” on page 56](#)
- [Section 6.1.3, “Comparison of RAID Performance,” on page 57](#)
- [Section 6.1.4, “Comparison of Disk Fault Tolerance,” on page 58](#)
- [Section 6.1.5, “Configuration Options for RAIDs,” on page 58](#)
- [Section 6.1.6, “Guidelines for Member Disks,” on page 58](#)
- [Section 6.1.7, “RAID 5 Algorithms for Distributing Stripes and Parity,” on page 59](#)
- [Section 6.1.8, “Device Mapper,” on page 61](#)
- [Section 6.1.9, “EVMS and the Multi-Disk \(md\) Driver,” on page 61](#)

6.1.1 What Is a Software RAID?

A RAID combines multiple devices into a multi-disk array to provide resiliency in the storage device and to improve storage capacity and I/O performance. If a disk fails, some RAID levels keep data available in a degraded mode until the failed disk can be replaced and its content reconstructed.

A software RAID provides the same high availability that you find in a hardware RAID. The key operational differences are described in the following table:

Table 6-1 *Comparison of Software and Hardware RAIDs*

Feature	Linux Software RAID	Hardware RAID
RAID function	Multi-disk (<code>md</code>) driver	RAID controller on the disk array
RAID processing	In the host server's processor	RAID controller on the disk array

Feature	Linux Software RAID	Hardware RAID
RAID levels	Linux supports levels 0, 1, 4, 5, and 10, depending on the configuration tool. Level 10 is available only as a non-nested algorithm in <code>mdadm</code> .	Varies by vendor
Member disks	Any disk	Same disk array

Nested RAID (also referred to as hybrid RAID by the Storage Networking Industry Association) are multi-level RAID built from basic RAID levels. In general, nested RAID devices for Linux are not supported because of how the `md` driver stores key information. The RAID information for the second RAID overwrites the first. If you attempt to nest the basic RAID, it appears to work, but the configuration is lost upon reboot.

The `mdadm` management tool supports only one complex RAID level, called RAID10. However, the RAID10 option is not formed by nesting basic RAID. It is similar to a RAID 1+0 but with striping integrated into the array. Data is not mirrored to a second disk. Instead, the algorithm uses dual-rotating striping to disks in the array. The MD plug-in to EVMS does not support the `md` driver's RAID10 option. To learn more about the RAID10 option for `mdadm`, see the `mdadm` man page.

NOTE: The Novell® Storage Services (NSS) file system for Linux supports software RAID 0, 1, 5, and nested RAID 10 and 15. NSS interfaces with EVMS APIs to manage devices on Linux, but also adds code that allows nesting on Linux. NSS is available in Novell Open Enterprise Server. For information, visit the [Novell Open Enterprise Server Web site \(http://www.novell.com/products/openenterpriseserver/\)](http://www.novell.com/products/openenterpriseserver/).

6.1.2 Overview of RAID Levels

The following table describes the advantages and disadvantages of the RAID levels supported by EVMS. The description assumes that the member devices reside on different disks and that each disk has its own dedicated I/O capability.

Table 6-2 RAID Levels Supported by EVMS

RAID Level	Description	Performance and Fault Tolerance
0	Stripes data using a round-robin method to distribute data over the RAID's member devices.	Improves disk I/O performance for both reads and writes. Actual performance depends on the stripe size, the actual data, and the application. Does not provide disk fault tolerance and data redundancy. Any disk failure causes all data in the RAID to be lost.

RAID Level	Description	Performance and Fault Tolerance
1	Mirrors data by copying blocks of one disk to another and keeping them in continuous synchronization. If disks are different sizes, the smaller disk determines the size of the RAID.	<p>Improves disk reads by making multiple copies of data available via different I/O paths. The write performance is about the same as for a single disk because a copy of the data must be written to each of the disks in the mirror.</p> <p>Provides 100% data redundancy. If one disk fails then the data remains available on its mirror, and processing continues.</p>
4	Stripes data and records parity to a dedicated disk. If disks are different sizes, the smaller disk determines the size of the RAID.	<p>Improves disk I/O performance for both reads and writes. Write performance is slightly less than for RAID 0, because parity must be calculated and written. Read performance is less than for mirrored data on a RAID 1.</p> <p>This type of striping is seldom used because if the parity disk is lost, the parity data cannot be reconstructed. The parity disk can become a bottleneck for I/O.</p>
5	Stripes data and distributes parity in a round-robin fashion across all disks. If disks are different sizes, the smaller disk determines the size of the RAID.	<p>Improves disk I/O performance for reads and writes. Write performance is less than for RAID 0, because parity must be calculated and written. Read performance is less than for mirrored data on a RAID 1. Actual performance depends on the stripe size, the actual data, and the application.</p> <p>Provides disk fault tolerance. If a disk fails, performance is degraded while the RAID uses the parity to reconstruct data for the replacement disk. Provides slightly less data redundancy than mirroring because it uses parity to reconstruct the data.</p>

6.1.3 Comparison of RAID Performance

The following table compares the read and write performance for RAID devices.

Table 6-3 *Read and Write Performance for RAIDs*

Raid Level	Read Performance	Write Performance
0	Faster than for a single disk	Faster than for a single disk and other RAIDs.
1	Faster than for a single disk, increasing as more mirrors are added	Slower than for a single disk, declining as more mirrors are added.
4	Faster than for a single disk. Slower than a RAID 0 because one disk is used for parity	Faster than for a single disk. Slower than a RAID 0 because of writes for parity. Slower than a RAID 5 because of possible bottlenecks for writes of parity to the same disk.
5	Faster than for a single disk; comparable to a RAID 0	Faster than a single disk. Slower than a RAID 0 because of writes for parity.

6.1.4 Comparison of Disk Fault Tolerance

The following table compares the disk fault tolerance for RAID devices.

Table 6-4 *Fault Tolerance for RAIDs*

Raid Level	Number of Disk Failures Tolerated	Data Redundancy
0	None	No
1	Number of disks minus 1	100% redundancy for each mirror
4	1, but not the parity disk	Dedicated parity disk to reconstruct data
5	1	Distributed parity to reconstruct data

6.1.5 Configuration Options for RAIDs

Table 6-5 *Configuration Options*

Option	Description
Spare Disk	<p>For RAIDs 1, 4, and 5, you can optionally specify a device, segment, or region to use as the replacement for a failed disk (the member device, segment, or region). On failure, the spare disk automatically replaces the failed disk, then reconstructs the data.</p> <p>However, if the parity disk fails on a RAID 4, parity cannot be reconstructed.</p>
Chunk Size (KB)	<p>For RAIDs 0, 4, or 5, specify the stripe size in KB.</p> <p>Consider the intended use of the RAID, such as the file system block size, the applications used, and the actual data (file sizes and typical reads and writes). A typical write size for large files is 128 KB.</p> <p>Default: 32 KB</p> <p>Range: 4 KB to 4096 KB, in powers of 2.</p>
RAID Level	<p>If you selected <i>MD RAID 4/5 Region Manager</i>, specify <i>RAID 4</i> or <i>RAID 5</i> (default).</p>
RAID Algorithm	<p>For RAID 5, specify one of the following algorithms to use for striping and distributing parity on the disk.</p> <ul style="list-style-type: none">• Left Asymmetric• Left Symmetric (Default, fastest performance for large reads)• Right Asymmetric• Right Symmetric

6.1.6 Guidelines for Member Disks

For efficient use of space and performance, the disks you use to create the RAID should have the same storage capacity. Typically, if member disks are not of identical storage capacity, then each

member of the RAID uses only an amount of space equal to the capacity of the smallest member disk.

The commonly used MS-DOS partition table format cannot be used on devices larger than 2 TB. The root (/) and /boot directories must be located on devices that are 2 TB in size or less. Because the md driver works with MS-DOS partition table format, the maximum size disk that can be a member of an md software RAID set is 2 TB.

When the member disks of the RAID are combined, the total size of the RAID device should not exceed the maximum supported device size that can be recognized by the file systems you plan to use. For example, a 32-bit system with Large File System Support limits file sizes to 2 TB. Some file systems use LVM to combine space from multiple disks to create devices that are larger than 2 TB. For information about the kernel limits for file system devices, see the [SUSE Linux Kernel Limits document \(http://www.novell.com/products/linuxenterpriseserver/kernel_limits.html\)](http://www.novell.com/products/linuxenterpriseserver/kernel_limits.html).

In general, each storage object included in the RAID should be from a different physical disk to maximize I/O performance and to achieve disk fault tolerance where supported by the RAID level you use. In addition, they should be of the same type (disks, segments, or regions).

Using member disks of differing speeds might introduce a bottleneck during periods of demanding I/O. The best performance can be achieved by using the same brand and models of disks and controllers in your hardware solution. If they are different, you should try to match disks and controllers with similar technologies, performance, and capacity. Use a low number of drives on each controller to maximize throughput.

IMPORTANT: As with any hardware solution, using the same brand and model introduces the risk of concurrent failures over the life of the product, so plan maintenance accordingly.

The following table provides recommendations for the minimum and maximum number of storage objects to use when creating a software RAID:

Table 6-6 *Recommended Number of Storage Objects to Use in the Software RAID*

RAID Type	Minimum Number of Storage Objects	Recommended Maximum Number of Storage Objects
RAID 0 (striping)	2	8
RAID 1 (mirroring)	2	4
RAID 4 (striping with dedicated parity)	3	8
RAID 5 (striping with distributed parity)	3	8

Connection fault tolerance can be achieved by having multiple connection paths to each storage object in the RAID. For more information about configuring multipath I/O support for your software RAID, see [Section 5.8, “Configuring Multipath I/O for Devices,” on page 49](#).

6.1.7 RAID 5 Algorithms for Distributing Stripes and Parity

RAID 5 uses an algorithm to determine the layout of stripes and parity. The following table describes the algorithms.

Table 6-7 RAID 5 Algorithms

Algorithm	EVMS Type	Description
Left Asymmetric	1	<p>Stripes are written in a round-robin fashion from the first to last member segment. The parity's position in the striping sequence moves in a round-robin fashion from last to first. For example:</p> <pre> sda1 sdb1 sdc1 sde1 0 1 2 p 3 4 p 5 6 p 7 8 p 9 10 11 12 13 14 p </pre>
Left Symmetric	2	<p>This is the default setting and is considered the fastest method for large reads.</p> <p>Stripes wrap to follow the parity. The parity's position in the striping sequence moves in a round-robin fashion from last to first. For example:</p> <pre> sda1 sdb1 sdc1 sde1 0 1 2 p 4 5 p 3 8 p 6 7 p 9 10 11 12 13 14 p </pre>
Right Asymmetric	3	<p>Stripes are written in a round-robin fashion from the first to last member segment. The parity's position in the striping sequence moves in a round-robin fashion from first to last. For example:</p> <pre> sda1 sdb1 sdc1 sde1 p 0 1 2 3 p 4 5 6 7 p 8 9 10 11 p p 12 13 14 </pre>
Right Symmetric	4	<p>Stripes wrap to follow the parity. The parity's position in the striping sequence moves in a round-robin fashion from first to last. For example:</p> <pre> sda1 sdb1 sdc1 sde1 p 0 1 2 5 p 3 4 7 8 p 6 9 10 11 p p 12 13 14 </pre>

For information about the layout of stripes and parity with each of these algorithms, see [Linux RAID-5 Algorithms \(http://www.accs.com/p_and_p/RAID/LinuxRAID.html\)](http://www.accs.com/p_and_p/RAID/LinuxRAID.html).

6.1.8 Device Mapper

The Device Mapper is not started by default in the rescue system.

- 1 Open a terminal console, then log in as the `root` user or equivalent.
- 2 Start the Device Mapper by entering the following at the terminal console prompt:

```
/etc/init.d/boot.device-mapper start
```

6.1.9 EVMS and the Multi-Disk (md) Driver

On Linux, the Multi-Disk (`md`) driver supports the following features for multiple devices:

- **Software RAIDs:** RAID 0 (striping), 1 (mirror), 4 (striping with dedicated parity), and 5 (striping with distributed parity).
- **Multipath I/O:** Connection fault tolerance and load balancing for connections between the server and disks where multiple paths are available.
- **Snapshots:** Snapshots of a file system at a particular point in time, even while the system is active, thereby allowing a consistent backup.
- **Linear RAID:** A linear concatenation of discontinuous areas of free space from the same or multiple storage devices. Areas can be of different sizes.

The MD plug-in to EVMS allows you to manage all of these MD features as “regions” with the Regions Manager.

6.2 Creating and Configuring a Software RAID

- 1 Open a terminal console, then log in as the `root` user or equivalent.
- 2 Start the EVMS GUI by entering the following at the terminal console prompt:

```
evmsgui
```
- 3 If the disks have not been initialized, initialize them by adding the DOS Segment Manager now.

The following instructions assume you are initializing new disks. For information about initializing an existing disk or a disk moved from another system, see [Section 4.2, “Initializing Disks,” on page 36](#).

Repeat the following steps for each disk that you want to initialize:

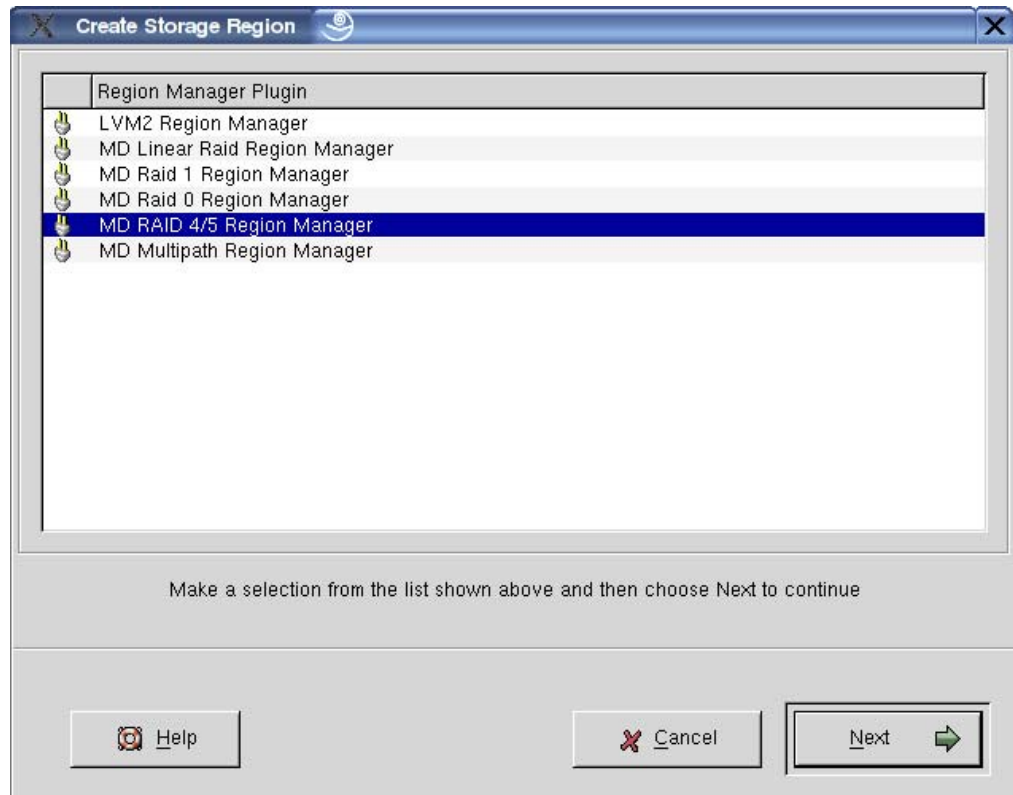
- 3a** Select *Actions > Add > Segment Manager to Storage Object*.
 - 3b** From the list, select the *DOS Segment Manager*, then click *Next*.
 - 3c** Select the device, then click *Add* to initialize it.
- 4 If segments have not been created on the disks, create a segment on each disk that you plan to use in the RAID.

For x86 platforms, this step is optional if you treat the entire disk as one segment.

For IA-64 platforms, this step is necessary to make the *RAID 4/5* option available in the Regions Manager.


For information about creating segments, see [Section 4.4, “Creating Disk Segments \(or Partitions\),” on page 38](#).

- 4a Select *Action > Create > Segment* to open the *DOS Segment Manager*.
- 4b Select the free space segment you want to use.
- 4c Specify the amount of space to use for the segment.
- 4d Specify the segment options, then click *Create*.
- 5 Create and configure a software RAID Device.
 - 5a Select *Action > Create > Region* to open the Create Storage Region dialog box.




- 5b Specify the type of software RAID you want to create by selecting one of the following Region Managers, then click *Next*.
 - *MD RAID 0 Region Manager*
 - *MD RAID 1 Region Manager*
 - *MD RAID 4/5 Region Manager*
- 5c From the Storage Objects listed, select the ones to use for the RAID device.

IMPORTANT: The order of the objects in the RAID is implied by their order in the list.



- 5d** Specify values for *Configuration Options* by changing the following default settings as desired.
- For RAID 1, 4, or 5, optionally specify a device to use as the spare disk for the RAID. The default is none.
 - For RAID 0, 4, or 5, specify the chunk (stripe) size in KB. The default is 32 KB.
 - For RAID 4/5, specify *RAID 4* or *RAID 5* (default).
 - For RAID 5, specify the algorithm to use for striping and parity. The default is *Left Symmetric*.

For information about these settings, see “Configuration Options for RAID5” on page 58.



- 5e** Click *Create* to create the RAID device under the `/dev/evms/md` directory.
The device is given a name such as `md0`, so its EVMS mount location is `/dev/evms/md/md0`.
- 6** Specify a human-readable label for the device.
 - 6a** Select *Action > Create > EVMS Volume or Compatible Volume*.
 - 6b** Select the device that you created in **Step 5**.
 - 6c** Specify a name for the device.
Use standard ASCII characters and naming conventions. Spaces are allowed.
 - 6d** Click *Done*.
- 7** Create a file system on the RAID device you created.
 - 7a** Select *Action > File System > Make* to view a list of file system modules.
 - 7b** Select the type of file system you want to create, such as the following:
 - *ReiserFS File System Module*
 - *Ext2/3FS File System Module*
 - 7c** Select the RAID device you created in **Step 5**, such as `/dev/evms/md/md0`.
 - 7d** Specify a name to use as the *Volume Label*, then click *Make*.
The name must not contain space or it will fail to mount later.
 - 7e** Click *Save* to create the file system.
- 8** Mount the RAID device.

- 8a** Select *Action > File System > Mount*.
- 8b** Select the RAID device you created in [Step 5](#), such as `/dev/evms/md/md0`.
- 8c** Specify the location where you want to mount the device, such as `/home`.
- 8d** Click *Mount*.
- 9** Enable `boot.evms` to activate EVMS automatically at reboot.
 - 9a** In YaST, select *System > System Services (Run Level)*.
 - 9b** Select *Expert Mode*.
 - 9c** Select *Boot.evms*.
 - 9d** Select *Set/Reset*.
 - 9e** Select *Enable the Service*.
- 10** Edit the `/etc/fstab` file to automount the RAID mount point created in [Step 8c](#), or you can mount the device manually from `evmsgui`.

6.3 Expanding a RAID

This section explains how to expand a RAID by adding segments to it.

IMPORTANT: Before you can expand the size of a RAID device, you must deactivate it.

- [Section 6.3.1, “Adding Mirrors to a RAID 1 Device,” on page 65](#)
- [Section 6.3.2, “Adding Segments to a RAID 4 or 5,” on page 66](#)

6.3.1 Adding Mirrors to a RAID 1 Device

In a RAID 1 device, each member segment contains its own copy of all of the data stored in the RAID. You can add a mirror to the RAID to increase redundancy. The segment must be at least the same size as the smallest member segment in the existing RAID 1 device. Any excess space in the segment is not used. Ideally, all member segments of a RAID 1 device are the same size.

Adding an Available Segment as the New Mirror

- 1** Deactivate the RAID 1 device.
- 2** Use the Add Active (`addactive` plug-in) function.
- 3** From the list of available segments, select one that is the same size or larger than the smallest existing member of the RAID device.
- 4** Reactivate the RAID device.

Activating A Spare Disk as the New Mirror

- 1** If you have not set up a spare disk, do it now.
 For information, see [Section 6.4, “Adding or Removing a Spare Disk,” on page 66](#).
- 2** Use the Activate Spare (`activatespare` plug-in) function to add it to the RAID 1 device as a new mirror.

6.3.2 Adding Segments to a RAID 4 or 5

If the RAID region is clean and operating normally, the kernel driver adds the new object as a regular spare, and it acts as a hot standby for future failures. If the RAID region is currently degraded, the kernel driver immediately activates the new spare object and begin synchronizing the data and parity information.

6.4 Adding or Removing a Spare Disk

The MD driver allows you to optionally designate a spare disk (device, segment, or region) for RAID 1, 4, and 5 devices. You can assign a spare disk when you create the RAID or at any time thereafter. The RAID can be active and in use when you add or remove the spare. The spare is activated for the RAID only on disk failure.

- [Section 6.4.1, “Do I Need a Spare Disk?” on page 66](#)
- [Section 6.4.2, “Adding a Spare Disk When You Create the RAID,” on page 66](#)
- [Section 6.4.3, “Adding a Spare Disk to an Existing RAID,” on page 66](#)
- [Section 6.4.4, “Removing a Spare Disk from a RAID,” on page 67](#)

6.4.1 Do I Need a Spare Disk?

The advantage of specifying a spare disk for a RAID is that the system monitors the failure and begins recovery without human interaction. The disadvantage is that the space on the spare disk is not available until it is activated by a failed RAID.

As noted in [“Overview of RAID Levels” on page 56](#), RAID 1, 4, and 5 can tolerate at least one disk failure. Any given RAID can have one spare disk designated for it, but the spare itself can serve as the designated spare for one RAID, for multiple RAID 1s, or for all arrays. The spare disk is a hot standby until it is needed. It is not an active member of any RAID 1s where it is assigned as the spare disk until it is activated for that purpose.

If a spare disk is defined for the RAID, the RAID automatically deactivates the failed disk and activates the spare disk on disk failure. The MD driver then begins synchronizing mirrored data for a RAID 1 or reconstructing the missing data and parity information for RAID 4 and 5. The I/O performance remains in a degraded state until the failed disk's data is fully remirrored or reconstructed.

6.4.2 Adding a Spare Disk When You Create the RAID

When you create a RAID 1, 4, or 5 in EVMS, specify the *Spare Disk* in the *Configuration Options* dialog box. You can browse to select the available device, segment, or region that you want to make the RAID's spare disk. For information, see [Step 5d in Section 6.2, “Creating and Configuring a Software RAID,” on page 61](#).

6.4.3 Adding a Spare Disk to an Existing RAID

The RAID 1, 4, or 5 device can be active and in use when you add a spare disk to it. If the RAID is operating normally, the specified disk is added as a spare and it acts as a hot standby for future failures. If the RAID is currently degraded because of a failed disk, the specified disk is added as a

spare disk, then it is automatically activated as a replacement disk for the failed disk, and it begins synchronizing the data and parity information.

- 1 Prepare a disk, segment, or region to use as the replacement disk, just as you did for the member devices of the RAID device.
- 2 In EVMS, select the *Actions > Add > Spare Disk to a Region* (the `addspare` plug-in for the EVMS GUI).
- 3 Select the RAID device you want to manage from the list of Regions, then click *Next*.
- 4 Select the device to use as the spare disk.
- 5 Click *Add*.

6.4.4 Removing a Spare Disk from a RAID

The RAID 1, 4, or 5 device can be active and in use when you remove its spare disk.

- 1 In EVMS, select the *Actions > Remove > Spare Disk from a Region* (the `remspare` plug-in for the EVMS GUI).
- 2 Select the RAID device you want to manage from the list of Regions, then click *Next*.
- 3 Select the spare disk.
- 4 Click *Remove*.

6.5 Removing a Failed Disk

RAIDs 1, 4, and 5 can survive a disk failure. A RAID 1 device survives if all but one mirrored array fails. Its read performance is degraded without the multiple data sources available, but its write performance might actually improve when it does not write to the failed mirrors. During the synchronization of the replacement disk, write and read performance are both degraded. A RAID 5 can survive a single disk failure at a time. A RAID 4 can survive a single disk failure at a time if the disk is not the parity disk.

Disks can fail for many reasons such as the following:

- Disk crash
- Disk pulled from the system
- Drive cable removed or loose
- I/O errors

When a disk fails, the RAID removes the failed disk from membership in the RAID, and operates in a degraded mode until the failed disk is replaced by a spare. Degraded mode is resolved for a single disk failure in one of the following ways:

- **Spare Exists:** If the RAID has been assigned a spare disk, the MD driver automatically activates the spare disk as a member of the RAID, then the RAID begins synchronizing (RAID 1) or reconstructing (RAID 4 or 5) the missing data.
- **No Spare Exists:** If the RAID does not have a spare disk, the RAID operates in degraded mode until you configure and add a spare. When you add the spare, the MD drive detects the RAID's degraded mode, automatically activates the spare as a member of the RAID, then begins synchronizing (RAID 1) or reconstructing (RAID 4 or 5) the missing data.

For information about adding a spare disk to a RAID, see [Section 6.4, “Adding or Removing a Spare Disk,” on page 66](#).

If you pull a disk or if it is totally unusable, EVMS no longer recognizes the failed disk as part of the RAID.

If you correct the problem that caused the disk to fail and leave the failed disk on the system, EVMS recognizes the disk as a faulty object in the RAID. The RAID does not use the disk because it is no longer synchronized with the RAID. EVMS does not make the device available for other use until you remove it from the RAID. Afterwards, the disk appears in the *Available-Objects* list in the EVMS GUI, where it can be used for any purpose.

The RAID 1, 4, or 5 device can be active and in use when you remove its faulty object.

- 1 In EVMS, select the *Actions > Remove > Faulty Object from a Region* (the `remfaulty` plug-in in the EVMS GUI).
- 2 Select the RAID device you want to manage from the list of Regions, then click *Next*.
- 3 Select the failed disk.
- 4 Click *Remove*.

6.6 Replacing a Failed Disk

RAIDs 1, 4, or 5 can tolerate a failed disk. You can either keep a spare assigned to the RAID as a hot standby to use as a replacement, or assign a spare disk to it as needed.

On failure, the `md` driver checks if a spare disk is assigned to the RAID. If a spare is available, it automatically removes the failed disk, replaces it with the spare disk, then begins to synchronize the data (for RAID 1) or reconstruct the data from parity (for RAID 4 or 5). You can remove the failed disk at any time after it has been replaced with the spare disk.

If a spare is not available, the RAID operates in degraded mode until you replace the drive.

- 1 Prepare the disk as needed to match the other members of the RAID.
- 2 Add the disk as a spare for the RAID.

For information about adding a spare disk, see [Section 6.4, “Adding or Removing a Spare Disk,” on page 66](#).

The `md` driver automatically begins the replacement and reconstruction or synchronization process.

- 3 Monitor the status of the RAID to verify the process has begun.

For information about how monitor RAID status, see [Section 6.7, “Monitoring Status for a RAID,” on page 68](#).

- 4 Remove the failed disk.

6.7 Monitoring Status for a RAID

- [Section 6.7.1, “Monitoring Status with EVMSGUI,” on page 69](#)
- [Section 6.7.2, “Monitoring Status with `/proc/mdstat`,” on page 69](#)
- [Section 6.7.3, “Monitoring Status with `mdadm`,” on page 70](#)
- [Section 6.7.4, “Monitoring a Remirror or Reconstruction,” on page 71](#)

6.7.1 Monitoring Status with EVMSGUI

The *Regions* tab in EVMS GUI (`evmsgui`) reports any software RAID devices that are defined and whether they are currently active.

6.7.2 Monitoring Status with `/proc/mdstat`

A summary of RAID and status information (active/not active) is available in the `/proc/mdstat` file.

- 1 Open a terminal console, then log in as the `root` user or equivalent.
- 2 View the `/proc/mdstat` file by entering the following at the console prompt:

```
cat /proc/mdstat
```

- 3 Evaluate the information.

The following table shows an example output and how to interpret the information.

Status Information	Description	Interpretation
Personalities : [raid5] [raid4]	List of the RAIDs on the server by RAID label.	You have two RAIDs defined with labels of <code>raid5</code> and <code>raid4</code> .
md0 : active raid5 sdg1[0] sdk1[4] sdj1[3] sdi1[2]	<device> : <active not active> <RAID label you specified> < storage object> [RAID order]	The RAID is active and mounted at <code>/dev/evms/md/md0</code> . The RAID label is <code>raid5</code> . The active segments are <code>sdg1</code> , <code>sdi1</code> , <code>sdj1</code> , and <code>sdk1</code> , as ordered in the RAID. The RAID numbering of 0 to 4 indicates that the RAID has 5 segments, and the second segment [1] is missing from the list. Based on the segment names, the missing segment is <code>sdh1</code> .
35535360 blocks level 5, 128k chunk, algorithm 2 [5/4] [U_UUU]	<number of blocks> blocks level < 0 1 4 5 > <stripe size in KB> chunk algorithm <1 2 3 4 > [number of devices/number of working devices] [U-UUU]	If the block size on the server is 4 KB, the total size of the RAID (including parity) is 142 GB, with a data capacity of 113.7 GB. The stripe size is 128 KB. The RAID is using left symmetric. algorithm <1 2 3 4 > [number of devices/number of working devices] [U-UUU]
unused devices: <none>	All segments in the RAID are in use.	There are no spare devices available on the server.

6.7.3 Monitoring Status with mdadm

To view the RAID status with the `mdadm` command, enter the following at a terminal prompt:

```
mdadm -D /dev/mdx
```

Replace *mdx* with the RAID device number.

Example 1: A Disk Fails

In the following example, only 4 of the 5 devices in the RAID are active (Raid Devices : 5, Total Devices : 4). When it was created, the member disks in the device were numbered 0 to 5 and are ordered according to their alphabetic appearance in the list where they were chosen, such as `/dev/sdg1`, `/dev/sdh1`, `/dev/sdi1`, `/dev/sdj1`, and `/dev/sdk1`. From the pattern of filenames of the other devices, you determine that the device that was removed was named `/dev/sdh1`.

```
/dev/md0:
  Version : 00.90.03
  Creation Time : Sun Apr 16 11:37:05 2006
    Raid Level : raid5
    Array Size : 35535360 (33.89 GiB 36.39 GB)
    Device Size : 8883840 (8.47 GiB 9.10 GB)
    Raid Devices : 5
    Total Devices : 4
Preferred Minor : 0
    Persistence : Superblock is persistent
    Update Time : Mon Apr 17 05:50:44 2006
      State : clean, degraded
    Active Devices : 4
    Working Devices : 4
    Failed Devices : 0
    Spare Devices : 0

    Layout : left-symmetric
    Chunk Size : 128K
    UUID : 2e686e87:1eb36d02:d3914df8:db197afe
    Events : 0.189

Number   Major   Minor   RaidDevice State   /dev/sdg1
    0         8       97         0   active sync
    1         0        0         1   removed
    2         8      129         2   active sync
    3         8      145         3   active sync
    4         8      161         4   active sync
```

Example 2: Spare Disk Replaces the Failed Disk

In the following `mdadm` report, only 4 of the 5 disks are active and in good condition (Active Devices : 4, Working Devices : 5). The failed disk was automatically detected and removed from the RAID (Failed Devices: 0). The spare was activated as the replacement

disk, and has assumed the diskname of the failed disk (/dev/sdh1). The faulty object (the failed disk that was removed from the RAID) is not identified in the report. The RAID is running in degraded mode (State : clean, degraded, recovering). The data is being rebuilt (spare rebuilding /dev/sdh1), and the process is 3% complete (Rebuild Status : 3% complete).

```
mdadm -D /dev/md0
/dev/md0:
  Version : 00.90.03
  Creation Time : Sun Apr 16 11:37:05 2006
  Raid Level : raid5
  Array Size : 35535360 (33.89 GiB 36.39 GB)
  Device Size : 8883840 (8.47 GiB 9.10 GB)
  Raid Devices : 5
  Total Devices : 5
  Preferred Minor : 0
  Persistence : Superblock is persistent
  Update Time : Mon Apr 17 05:50:44 2006
  State : clean, degraded, recovering
  Active Devices : 4
  Working Devices : 5
  Failed Devices : 0
  Spare Devices : 1
  Layout : left-symmetric
  Chunk Size : 128K
  Rebuild Status : 3% complete
  UUID : 2e686e87:1eb36d02:d3914df8:db197afe
  Events : 0.189

Number   Major   Minor   RaidDevice State
  0         8       97         0      active sync   /dev/sdg1
  1         8      113         1      spare rebuilding /dev/sdh1
  2         8      129         2      active sync   /dev/sdi1
  3         8      145         3      active sync   /dev/sdj1
  4         8      161         4      active sync   /dev/sdk1
```

6.7.4 Monitoring a Remirror or Reconstruction

You can follow the progress of the synchronization or reconstruction process by examining the /proc/mdstat file.

You can control the speed of synchronization by setting parameters in the /proc/sys/dev/raid/speed_limit_min and /proc/sys/dev/raid/speed_limit_max files. To speed up the process, echo a larger number into the speed_limit_min file.

6.8 Deleting a Software RAID and Its Data

If you want to remove the prior multipath settings, deactivate the RAID, delete the data on the RAID, and release all resources used by the RAID, you can clear the md superblock as follows:

- 1 If you want to keep the data stored on the software RAID device, make sure to back up the data to alternate media, using your normal backup procedures. Make sure the backup is good before proceeding.
- 2 Open a terminal console prompt as the `root` user or equivalent. Use this console to enter the commands described in the remaining steps.
- 3 Unmount the software RAID device by entering

```
umount <raid-device>
```
- 4 Stop the RAID device and its member devices by entering

```
mdadm --stop <raid-device>
mdadm --stop <member-devices>
```

For more information about using `mdadm`, please see the `mdadm(8)` man page.
- 5 Delete all data on the disk by literally overwriting the entire device with zeroes. Enter

```
mdadm --misc --zero-superblock <member-devices>
```
- 6 You must now reinitialize the disks for other uses, just as you would when adding a new disk to your system.

Installing and Managing DRBD Services

7

This section describes how to install, configure, and manage a device-level software RAID 1 across a network using DRBD* (Distributed Deplicated Block Device) for Linux. (A standard software RAID 1 device mirrors partitions on devices.)

- [Section 7.1, “Understanding DRBD,” on page 73](#)
- [Section 7.2, “Installing DRDB Services,” on page 73](#)
- [Section 7.3, “Configuring the DRDB Service,” on page 74](#)

7.1 Understanding DRBD

DRBD allows you to create a mirror of devices that are located at two different sites across a network. When used with HeartBeat 2 (HB2), DRBD supports distributed high-availability Linux clusters.

For information about DRBD, see [DRBD \(http://linux-ha.org/DRBD\)](http://linux-ha.org/DRBD) by the Linux High-Availability Project.

For information about installing and configuring HeartBeat 2 for SUSE® Linux Enterprise Server 10, see the [HeartBeat 2 Installation and Setup Guide \(http://www.novell.com/documentation/sles10/hb2/data/hb2_config.html\)](http://www.novell.com/documentation/sles10/hb2/data/hb2_config.html) on the [Novell Documentation Web site for SUSE Linux Enterprise Server 10 \(http://www.novell.com/documentation/sles10\)](http://www.novell.com/documentation/sles10).

7.2 Installing DRDB Services

- 1 Install the High Availability (HA) pattern on both SUSE Linux Enterprise Servers in your networked cluster. Installing HA also installs the `drbd` program files.
 - 1a Log in as the `root` user or equivalent, then open YaST.
 - 1b Choose *Software > Software Management*.
 - 1c Change the filter to *Patterns*.
 - 1d Under *Base Technologies*, select *High Availability*.
 - 1e Click *Accept*.
- 2 Install the `drbd` kernel modules on both servers.
 - 2a Log in as the `root` user or equivalent, then open YaST.
 - 2b Choose *Software > Software Management*.
 - 2c Change the filter to *Search*.
 - 2d Type `drbd`, then click *Search*.
 - 2e Select all of the `drbd-kmp-*` packages.
 - 2f Click *Accept*.
 - 2g When prompted to restart to apply the kernel changes, reboot the system.

7.3 Configuring the DRDB Service

NOTE: The following procedure uses the server names node 1 and node 2, and the cluster resource name r0. It sets up node 1 as the primary node. Make sure to modify the instructions to use your own node and file names.

- 1 Open the `/etc/drbd.conf` file on the primary node (node 1) in a text editor, modify the following parameters in the `on hostname { }` sections, then save the file.

- `hostname`
- `device`
- `disk`
- `meta-disk`
- `address`
- `port`

All of these options are explained in the examples in the `/usr/share/doc/packages/drbd/drbd.conf` file.

- 2 Copy the `/etc/drbd.conf` file to the `/etc/drbd.conf` location on the secondary server (node 2).

- 3 Configure the DRDB service for node 1.

- 3a Open a terminal console for node 1, then log in as the `root` user or equivalent.

- 3b Initialize DRDB on node 1 by entering

```
modprobe drbd
```

- 3c Test the configuration file by running `drbdadm` with the `-d` option. Enter

```
drbdadm -d adjust r0
```

- 3d If the partitions and settings are correct, run `drbdadm` again without the `-d` option. Enter

```
drbdadm adjust r0
```

Make sure there are no errors before continuing.

- 4 Configure the DRDB service for node 2.

- 4a Open a terminal console for node 2, then log in as the `root` user or equivalent.

- 4b Initialize the DRDB service on node 2 by entering

```
modprobe drbd
```

- 4c Test the configuration file by running `drbdadm` with the `-d` option. Enter

```
drbdadm -d adjust r0
```

- 4d If the partitions and settings are correct, run `drbdadm` again without the `-d` option. Enter

```
drbdadm adjust r0
```

Make sure there are no errors before continuing.

- 5 Configure `node1` as the primary node by entering

```
drbdsetup /dev/drbd0 primary -do-what-I-say
```

- 6 Start the DRDB service on both systems by entering the following on each node:

```
service drbd start
```

- 7** Check the DRDB service status by entering the following on each node:

```
service drbd status
```

- 8** Format the DRDB device on the primary with a file system such as reiserfs. Any Linux file system can be used. Enter the following:

```
mkfs.reiserfs -f /dev/drbd0
```

Always use the /dev/drbd# name in the command, not the actual /dev/disk device name.

- 9** Test the DRDB service on node 1.

- 9a** Create a mount point on node 1, such as /r0mount, by entering

```
mkdir /r0mount
```

- 9b** Mount the drbd device by entering

```
mount -o rw /dev/drbd0 /r0mount
```

- 9c** Create a file from the primary node by entering

```
touch /r0mount/from_node1
```

- 10** Test the DRDB service on node 2.

- 10a** Unmount the disk on node 1 by typing the following command on node 1:

```
umount /r0mount
```

- 10b** Downgrade the DRDB service on node 1 by typing the following command on node 1:

```
drbdadm secondary r0
```

- 10c** On node 2, promote the DRDB service to primary by entering

```
drbdadm primary r0
```

- 10d** On node 2, check to see if node 2 is primary by entering

```
service drbd status
```

- 10e** On node 2, create a mount point such as /r0mount, by entering

```
mkdir /r0mount
```

- 10f** On node 2, mount the DRDB device by entering

```
mount -o rw /dev/drbd0 /r0mount
```

- 10g** Verify that the file you created on node 1 in [Step 9c](#) is viewable by entering

```
ls /r0mount
```

The /r0mount/from_node1 file should be listed.

- 11** If the service is working on both nodes, the DRDB setup is complete.

- 12** Set up node 1 as the primary again.

- 12a** Unmount the disk on node 2 by typing the following command on node 2:

```
umount /r0mount
```

- 12b** Downgrade the DRDB service on node 2 by typing the following command on node 2:

```
drbdadm secondary r0
```

- 12c** On node 1, promote the DRDB service to primary by entering

```
drbdadm primary r0
```

12d On node 1, check to see if node 1 is primary by entering

```
service drbd status
```

- 13** To get the service to automatically start and fail over if the server has a problem, you can set up DRDB as a high availability service with HeartBeat 2.

Troubleshooting EVMS Devices, RAIDs, and Volumes

8

This section describes workarounds for known issues for EVMS devices, software RAIDs, multipath I/O, and volumes.

- [Section 8.1, “EVMS Volumes Might Not Appear When Using iSCSI,” on page 77](#)

8.1 EVMS Volumes Might Not Appear When Using iSCSI

If you have installed and configured an iSCSI SAN, and have created and configured EVMS Disks or Volumes on that iSCSI SAN, your EVMS volumes might not be visible or accessible. This problem is caused by EVMS starting before the iSCSI service. iSCSI must be started and running before any disks/volumes on the iSCSI SAN can be accessed.

To resolve this problem, use the `chkconfig` command at the Linux server console of every server that is part of your iSCSI SAN to correct the order that iSCSI and EVMS are started.

- 1 At a terminal console prompt, enter

```
chkconfig boot.evms on
```

This ensures that EVMS and iSCSI start in the proper order each time your servers reboot.