Tru64 UNIX

Logical Storage Manager

Part Number: AA-RH9BC-TE

June 2001

Product Version: Tru64 UNIX Version 5.1A or higher

This manual describes how to configure and manage disk storage using the Logical Storage Manager (LSM) software. It includes information on LSM concepts and how to plan, set up, monitor, change, and troubleshoot an LSM configuration.

Compaq Computer Corporation
Houston, Texas

Contents

5 Managing LSM Objects

7 Error Messages

A The Storage Administrator (lsmsa)

Glossary

Index

Examples

Figures

Tables

About This Manual

This manual describes how to configure and manage disk storage using the Logical Storage Manager (LSM) software. It includes information on LSM concepts and how to plan, set up, monitor, change, and troubleshoot an LSM configuration.

Audience

This manual is intended for anyone who needs to configure and manage storage devices under LSM control. To use this manual you must be able to administer a system running the Compaq Tru64™ UNIX operating system software and its storage devices.

New and Changed Features

This LSM manual has been revised to:

- Include information on using LSM in a TruCluster™ environment and references to the *Cluster Administration* manual for further information not covered here
- Clarify terminology and definitions
- Include information on converting disk groups from previous software versions
- Describe enhancements to the Storage Administrator GUI for managing AdvFS

Organization

This manual contains the following chapters and appendices.

Chapter 1	Describes LSM features, terms, and concepts, and introduces the available interfaces to LSM.
Chapter 2	Provides worksheets to aid you in planning your LSM configuration.
Chapter 3	Describes how to upgrade a system with an existing LSM configuration or install LSM for the first time as part of a system upgrade or installation.

Chapter 4	Describes how to create LSM disk groups and volumes and how to encapsulate existing data or AdvFS domains into LSM volumes.
Chapter 5	Describes common management tasks for LSM objects.
Chapter 6	Describes how to identify and recover from problems and how to replace disks under LSM control.
Chapter 7	Describes LSM error messages and solutions.
Appendix A	Describes how to install and start the Storage Administrator GUI and how to manage LSM objects using this interface.
Appendix B	Describes how to track Storage Administrator activities, how to use the Storage Administrator, and how to customize the Storage Administrator GUI.
Appendix C	Describes how to manage LSM objects using the `voldiskadm` menu interface.
Appendix D	Describes how to start the Visual Administrator (`dxlsm`) interface, and describes its windows, icons, and mouse operations.
Appendix E	Describes how to manage LSM objects with the Visual Administrator (`dxlsm`) interface.
Glossary	

Related Documentation

The following operating system documents provide information related to LSM:

- *Installation Guide* describes how to install the LSM software.
- *Release Notes* describe LSM problems and solutions that might not be documented elsewhere.
- *System Administration* describes general storage administration.
- *System Configuration and Tuning* describes how to plan, configure, and tune storage devices.
- *AdvFS Administration* describes how to use the AdvFS software with LSM.
- *Cluster Administration* describes how to configure LSM in a TruCluster environment.

Icons on Tru64 UNIX Printed Manuals

The printed version of the Tru64 UNIX documentation uses letter icons on the spines of the manuals to help specific audiences quickly find the manuals that meet their needs. (You can order the printed documentation from Compaq.) The following list describes this convention:

G	Manuals for general users
S	Manuals for system and network administrators
P	Manuals for programmers
R	Manuals for reference page users

Some manuals in the documentation help meet the needs of several audiences. For example, the information in some system manuals is also used by programmers. Keep this in mind when searching for information on specific topics.

The *Documentation Overview* provides information on all of the manuals in the Tru64 UNIX documentation set.

Reader's Comments

Compaq welcomes any comments and suggestions you have on this and other Tru64 UNIX manuals.

You can send your comments in the following ways:

- Fax: 603-884-0120 Attn: UBPG Publications, ZKO3-3/Y32
- Internet electronic mail: `readers_comment@zk3.dec.com`

  A Reader's Comment form is located on your system in the following location:

  `/usr/doc/readers_comment.txt`

Please include the following information along with your comments:

- The full title of the manual and the order number. (The order number appears on the title page of printed and PDF versions of a manual.)
- The section numbers and page numbers of the information on which you are commenting.
- The version of Tru64 UNIX that you are using.
- If known, the type of processor that is running the Tru64 UNIX software.

The Tru64 UNIX Publications group cannot respond to system problems or technical support inquiries. Please address technical questions to your local system vendor or to the appropriate Compaq technical support office.

Information provided with the software media explains how to send problem reports to Compaq.

Conventions

This manual uses the following conventions:

#	A number sign represents the superuser prompt.
`%` **`cat`**	Boldface type in interactive examples indicates typed user input.
`file`	Italic (slanted) type indicates variable values, placeholders, and function argument names.
[|] { | }	In syntax definitions, brackets indicate items that are optional and braces indicate items that are required. Vertical bars separating items inside brackets or braces indicate that you choose one item from among those listed.
. . .	In syntax definitions, a horizontal ellipsis indicates that the preceding item can be repeated one or more times.
`cat`(1)	A cross-reference to a reference page includes the appropriate section number in parentheses. For example, `cat`(1) indicates that you can find information on the `cat` command in Section 1 of the reference pages.

1 Overview

The Logical Storage Manager (LSM) software is an optional integrated, host-based disk storage management application that allows you to manage storage devices without disrupting users or applications accessing data on those storage devices.

LSM uses Redundant Arrays of Independent Disks (RAID) technology to enable you to configure storage devices into a virtual pool of storage from which you create LSM volumes. You configure new file systems, databases, and applications, or encapsulate existing ones, to use an LSM volume instead of a disk partition.

The benefits of using an LSM volume instead of a disk partition include:

- Data loss protection

  You can configure LSM to protect against data loss by configuring LSM volumes in one of the following ways:

  – To store and maintain multiple copies (mirrors) of data on different storage devices. If a storage device fails, LSM continues operating using mirror data.

  – To store data and parity information on different storage devices. If a storage device fails, LSM uses the data on the remaining storage devices and the parity information to reconstruct the missing data on the failed storage device.

  In either case, data remains available without disrupting users or applications, shutting down the system, or backing up and restoring data.

  You can configure LSM to encapsulate the boot disk partitions into LSM volumes and then create mirrors of those volumes. By doing so, you create copies of the boot disk partitions from which the system can boot if the original boot disk fails.

- Maximized disk usage

  You can configure LSM to seamlessly join together storage devices to appear as a single storage device to users and applications.

- Performance improvements

  You can configure LSM to separate data into units of equal size, then read or write the data units on two or more storage devices. LSM simultaneously reads or writes the data units if the storage devices are on different SCSI buses.

- Data availability

  You can configure LSM in a TruCluster environment. TruCluster software makes AlphaServer™ systems appear as a single system on the network. The AlphaServer systems running the TruCluster software become members of the cluster and share resources and data storage. This sharing allows applications, such as LSM, to continue uninterrupted if the cluster member on which it was running fails.

This chapter introduces LSM features, concepts, and terminology. The `volintro`(8) reference page also provides information on LSM terms and commands.

1.1 LSM Object Hierarchy

LSM uses the following hierarchy of objects to organize storage:

- LSM disk—An object that represents a storage device that is initialized exclusively for use by LSM
- Disk Group—An object that represents a collection of LSM disks and subdisks for use by an LSM volume
- Subdisk—An object that represents a contiguous set of blocks on an LSM disk that LSM uses to write volume data
- Plex—An object that represents a subdisk or collection of subdisks to which LSM writes a copy of the volume data or log information
- Volume—An object that represents a hierarchy of LSM objects, including LSM disks, subdisks, and plexes in a disk group. Applications and file systems make read and write requests to the LSM volume.

The following sections describe LSM objects in more detail.

1.1.1 LSM Disk

An LSM disk is a Tru64 UNIX supported storage device, including disks, disk partitions, and hardware RAID sets, that you configure exclusively for use by LSM. LSM views the storage in the same way as the Tru64 UNIX operating system software views it. For example, if the operating system software considers a RAID set as a single storage device, so does LSM.

For more information on supported storage devices, see the *Tru64 UNIX Software Product Description (SPD)* web site at the following URL:

`http://www.tru64unix.compaq.com/docs/spds.html`

Note

LSM does not recognize and support disk clones (hardware disk copies of LSM disks).

Figure 1–1 shows a typical hardware configuration that LSM supports.

Figure 1–1: Typical LSM Hardware Configuration

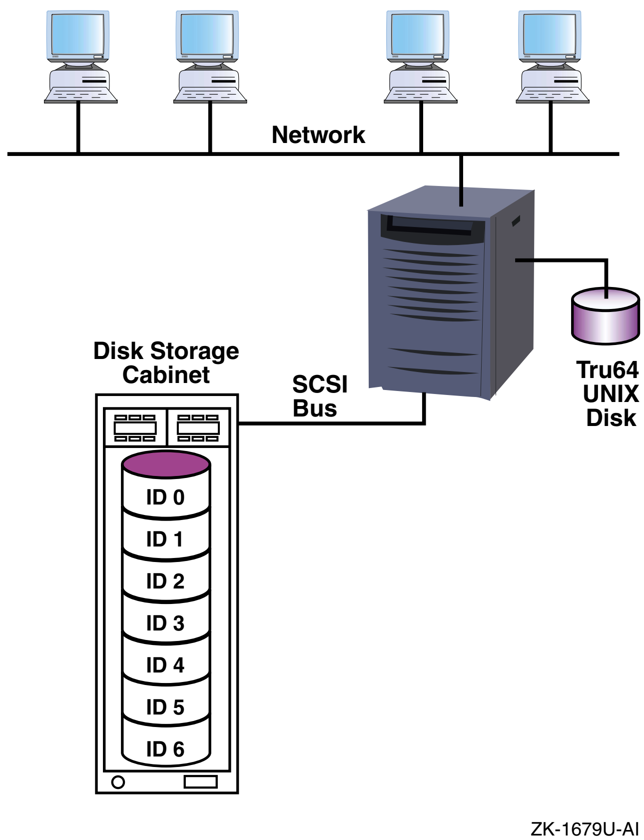

A storage device becomes an LSM disk when you initialize it for use by LSM. There are three types of LSM disks:

- A **sliced disk**, which initializes an entire disk for LSM use. This type of initialization organizes the storage into two regions on separate partitions—a large public region used for storing data and a private region for storing LSM internal metadata, such as LSM configuration information. The default size of the private region is 4096 blocks. Figure 1–2 shows a sliced disk.

Figure 1–2: LSM Sliced Disk

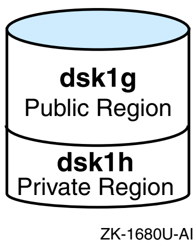

- A **simple disk**, which initializes a disk partition. This type of initialization organizes the storage into two regions on the same partition—a large public region used for storing data and a private region for storing LSM internal metadata, such as LSM configuration information. The default size of the private region is 4096 blocks. Figure 1–3 shows a simple disk.

  Figure 1–3: LSM Simple Disk

  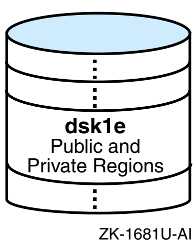

  Whenever possible, initialize the entire disk as a sliced disk instead of configuring individual disk partitions as simple disks. This ensures that the disk's storage is used efficiently and avoids using space for multiple private regions on the same disk.

- A **nopriv disk**, which initializes a disk or disk partition that contains data you want to encapsulate. This type of initialization creates only a public region for the data and no private region. Figure 1–4 shows a nopriv disk.

  Figure 1–4: LSM Nopriv Disk

  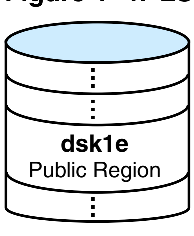

1.1.2 Disk Group

A disk group is an object that represents a grouping of LSM disks. LSM disks in a disk group share a common configuration database that identifies all the LSM objects in the disk group. LSM automatically creates and maintains copies of the configuration database in the private region of multiple LSM sliced or simple disks in each disk group.

LSM distributes these copies across all controllers for redundancy. If LSM disks in a disk group are located on the same controller, LSM distributes the copies across several disks. LSM automatically records changes to the LSM configuration and, if necessary, changes the number and location of copies of the configuration database for a disk group.

You cannot have a disk group of only LSM nopriv disks, because an LSM nopriv disk does not have a private region to store copies of the configuration database.

By default, the LSM software creates a default disk group called rootdg. The configuration database for rootdg contains information for itself and all other disk groups that you create.

An LSM volume can use disks only within the same disk group. You can create all of your volumes in the rootdg disk group, or you can create other disk groups. For example, if you dedicate disks to store financial data, you can create and assign those disks to a disk group called finance.

When you add an LSM disk to a disk group, LSM assigns it a disk media name. By default, the disk media name is the same as the disk access name, which the operating system software assigns to a storage device. For example, the disk media name and disk access name might be dsk1.

You do not have to use the default disk media name. You can assign a disk media name of up to 31 alphanumeric characters that cannot include spaces or the forward slash (/). For example, you could assign a disk media name of finance_data_disk.

LSM associates the disk media name with the operating system's disk access name. The disk media name provides insulation from operating system naming conventions. This allows LSM to find the device should you move it to a new location (for example, connect a disk to a different controller). However, LSM nopriv disks require more planning to move them to a different controller or a different system. See Section 5.2.7 for more information on moving a disk group containing nopriv disks to another system.

1.1.3 Subdisk

A subdisk is an object that represents a contiguous set of blocks in an LSM disk's public region that LSM uses to store data.

By default, LSM assigns a subdisk name using the LSM disk media name followed by a dash (–) and an ascending two-digit number beginning with 01. For example, dsk1–01 is the subdisk name on an LSM disk with a disk media name of dsk1.

You do not have to use the default subdisk name. You can assign a subdisk name of up to 31 alphanumeric characters that cannot include spaces or the forward slash (/). For example, you could assign a subdisk name of finance_disk01.

A subdisk can be:

- The entire public region. Figure 1–5 shows that the entire public region of an LSM disk was configured as a subdisk called dsk1–01:

Figure 1–5: Single Subdisk Using a Public Region

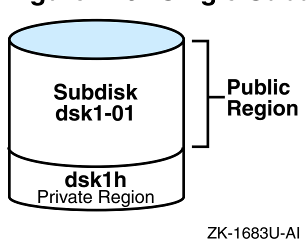

- A portion of the public region. Figure 1–6 shows a public region of an LSM disk that was configured as two subdisks called dsk2-01 and dsk2-02:

Figure 1–6: Multiple Subdisks Using a Public Region

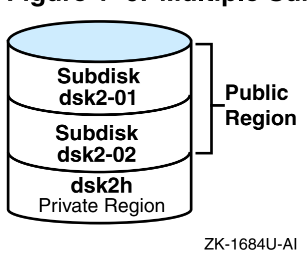

1.1.4 Plex

A plex is an object that represents a subdisk or collection of subdisks in the same disk group to which LSM writes a copy of volume data or log information. There are three types of plexes:

- Data plex

  A data plex contains volume data. There are three types of data plexes. The data plex that you choose depends on how you want LSM to store volume data on subdisks. The following lists the three types of data plexes:

  – **Concatenated data plex**

    In a concatenated data plex, LSM writes volume data in a linear manner. When the space in one subdisk has been written to, the remaining data goes to the next sequential subdisk in the plex. Section 1.1.4.1 explains this plex type in more detail.

  – **Striped data plex**

    In a striped data plex, LSM separates data into equal-sized units and writes the data units to each disk in the plex. This spreads the read-write operations evenly across the disks. Section 1.1.4.2 explains this plex type in more detail.

  – **RAID 5 data plex**

    In a RAID 5 data plex, LSM calculates a parity value for the data being written, then separates the data into equal-sized units and intersperses the data units and parity on each column in the plex. Section 1.1.4.3 explains this plex type in more detail.

- Log plex

  A log plex contains information about activity in a volume. In the event of a failure, LSM recovers only those areas of the volume identified in the log plex as being dirty (written to) at the time of the failure. There are two types of log plexes:

  – **Dirty Region Log (DRL) plex**

    In a DRL plex, LSM logs regions in a mirrored concatenated or striped data plex.

  – **RAID 5 log plex**

    In a RAID 5 log plex, LSM logs blocks being changed in a RAID 5 data plex and stores a temporary copy of the data and parity being written.

- Data and log plex (for compatibility with Version 4.0)

By default, LSM assigns a plex name using the volume name followed by a dash (–) and an ascending two-digit number beginning with 01. For example, volume1–01 is the name of a plex for a volume called volume1.

You do not have to use the default plex name. You can assign a plex name of up to 31 alphanumeric characters that cannot include spaces or the forward slash (/). For example, you could assign a plex media name of finance_plex01.

1.1.4.1 Concatenated Data Plex

In a concatenated data plex, LSM creates a contiguous address space on the subdisks and sequentially writes volume data in a linear manner. If LSM reaches the end of a subdisk while writing data, it continues to write data to the next subdisk as shown in Figure 1–7.

Figure 1–7: Concatenated Data Plex

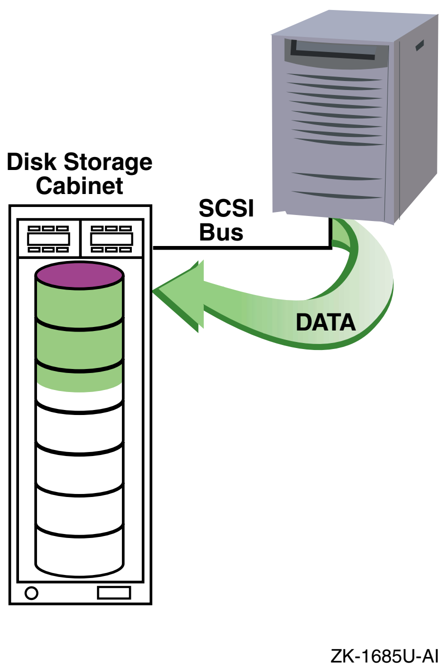

A single subdisk failure in a volume with one concatenated data plex will result in LSM volume failure. To prevent this type of failure, you can create multiple plexes (mirrors) on different disks. LSM continuously maintains the data in the mirrors. If a plex becomes unavailable because of a disk failure, the volume continues operating using another plex.

Using disks on different SCSI buses for mirror plexes speeds read requests, because data can be simultaneously read from multiple plexes.

By default, LSM creates a DRL plex when you create a mirrored volume. A DRL plex divides the data plexes into a set of consecutive regions and tracks regions that change due to I/O writes. When the system restarts after a failure, only the changed regions of the volume are recovered.

If you do not use a DRL plex and the system restarts after a failure, LSM must copy and resynchronize all the data to each plex to restore the plex consistency. Although this process occurs in the background and the volume is still available, it can be a lengthy procedure and can result in unnecessarily recovering data, thereby degrading system performance.

You can create up to 32 plexes, which can be any combination of data or DRL plexes.

Figure 1–8 shows a volume with mirrored concatenated data plexes.

Figure 1–8: Volume with Mirrored Concatenated Data Plexes

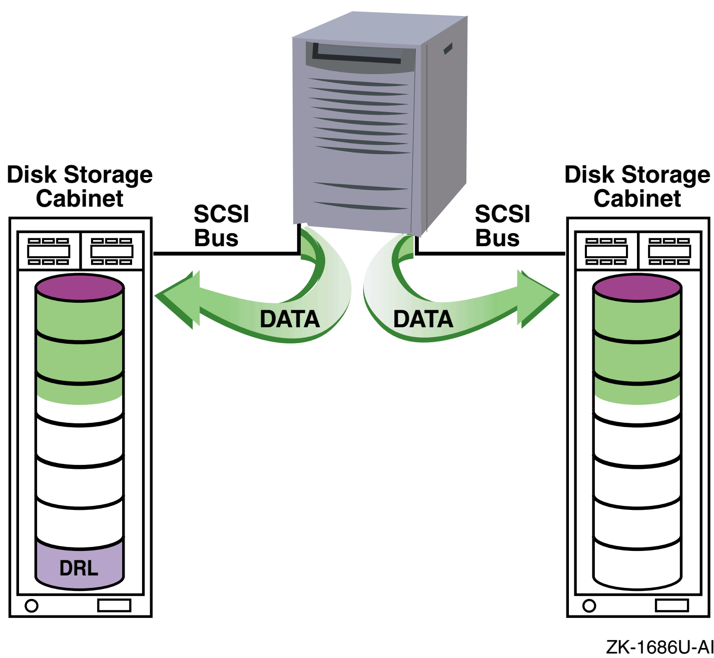

1.1.4.2 Striped Data Plex

In a striped data plex, LSM separates the data into units of equal size (64 KB by default) and writes the data units alternately on two or more columns of subdisks, creating a stripe of data across the columns. LSM can simultaneously write the data units if there are two or more units and the subdisks are on different SCSI buses.

Figure 1–9 shows how a write request of 384 KB of data is separated into six 64 KB units and written to three columns as two complete stripes.

Figure 1–9: Writing Data to a Striped Plex

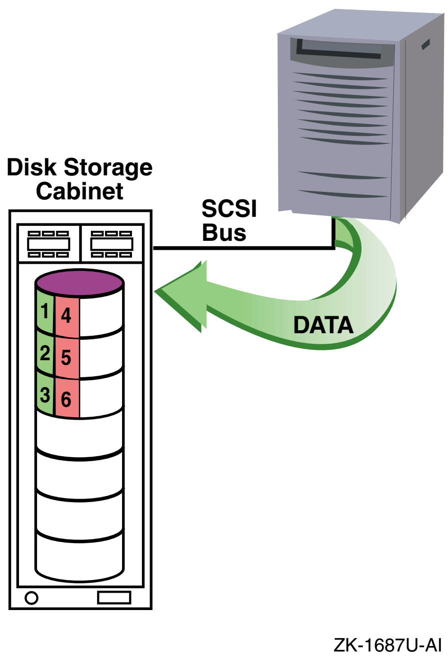

If a write request does not complete a stripe, then the first data unit of the next write request starts in the next column. For example, Figure 1–10 shows how 320 KB of data is separated into five 64 KB units and written to three columns. The first data unit of the next write request will start in the third column.

Figure 1–10: Incomplete Striped Data Plex

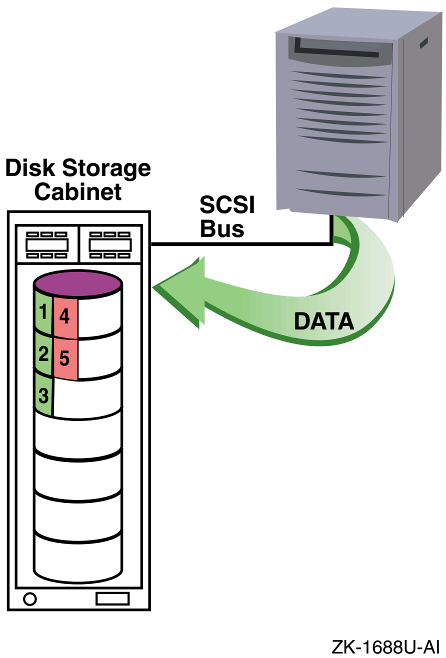

As in a concatenated data plex, a single disk failure in a volume with one striped data plex will result in volume failure. To prevent this type of failure, you can create multiple plexes (mirrors) on different disks. LSM continuously maintains the data in the mirrors. If a plex becomes unavailable because of a disk failure, the volume continues operating using another plex.

Using disks on different SCSI buses for mirror plexes speeds read requests, because data can be simultaneously read from multiple plexes.

By default, LSM creates a DRL plex when you create a mirrored volume. A DRL plex divides the data plexes into a set of consecutive regions and tracks regions that change due to I/O writes. When the system restarts after a failure, only the changed regions of the volume are recovered.

If you do not use a DRL plex and the system restarts after a failure, LSM must copy and resynchronize all the data to each plex to restore the plex consistency. Although this process occurs in the background and the volume is still available, it can be a lengthy procedure and can result in unnecessarily recovering data, thereby degrading system performance.

You can create up to 32 plexes, which can be any combination of data or DRL plexes.

Figure 1–11 shows a volume with mirrored striped data plexes.

Figure 1–11: Volume with Mirrored Striped Data Plexes

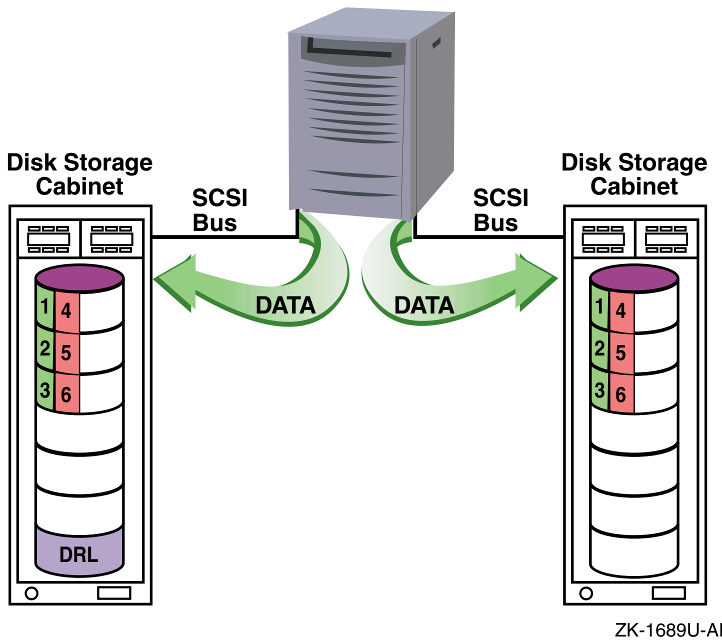

1.1.4.3 RAID 5 Data Plex

In a RAID 5 data plex, LSM calculates a parity value for each stripe of data, then separates the stripe of data and parity into units of equal size (16 KB by default) and writes the data and parity units on three or more columns of subdisks, creating a stripe of data across the columns. LSM can simultaneously write the data units if there are three or more units and the disks are on different SCSI buses. If a disk in one column fails, LSM continues operating using the data and parity information in the remaining columns to reconstruct the missing data.

In a RAID 5 data plex, LSM writes both data and parity across columns, writing the parity in a different column for each stripe of data. The first parity unit is located in the last column. Each successive parity unit is located in the next column, left-shifted one column from the previous parity unit location. If there are more stripes than columns, the parity unit placement begins again in the last column.

Figure 1–12 shows how data and parity information are written in a RAID 5 data plex with three columns.

Figure 1–12: Data and Parity Placement in a Three-Column RAID 5 Data Plex

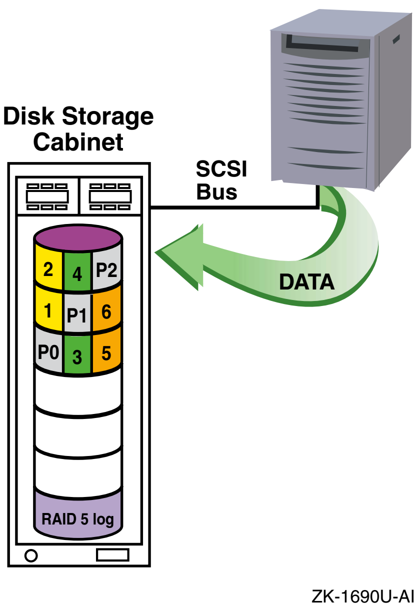

In Figure 1–12, the first stripe of data contains data units 1 and 2 and parity unit P0. The second stripe contains data units 3 and 4 and parity unit P1. The third stripe contains units 5 and 6 and parity unit P2.

By default, creating a RAID 5 volume creates a RAID 5 log plex. A RAID 5 log plex keeps track of data and parity blocks being changed due to I/O writes. When the system restarts after a failure, the write operations that did not complete before the failure are restarted.

Note

You cannot mirror a RAID 5 data plex.

The TruCluster software does not support RAID 5 volumes and data plexes.

1.1.5 LSM Volume

A volume is an object that represents a hierarchy of plexes, subdisks, and LSM disks in a disk group. Applications and file systems make read and write requests to the LSM volume. The LSM volume depends on the underlying LSM objects to satisfy the request.

An LSM volume can use storage from only one disk group.

A Note About Terminology

Volumes that use mirror plexes, whether those plexes are concatenated or striped, are often called *mirrored volumes*. If the plex layout is important, the volume might be described as a *concatenated and mirrored volume* or a *striped and mirrored volume*. Volumes that use a RAID plex are similarly called *RAID 5 volumes*. It is important to understand that these terms are a shorthand way of describing a volume that uses plexes of a certain type:

- A *mirrored volume* has two or more striped or concatenated plexes that each contain exact copies of the data. The volume is the container for all the copies (plexes).
- A *RAID 5 volume*, or a volume that uses a RAID 5 plex, can never be mirrored, because by definition and design, the volume contains only one data plex (and up to 31 log plexes). The data plex provides redundancy for the volume in the form of the parity value for each stripe of data. The log plex provides the fast-recovery mechanism for the volume by tracking the regions that change, along with a copy of the data and parity for a predefined number of writes.

As with most storage devices, an LSM volume has a block device interface and a character device interface.

- A volume's block device interface is located in the `/dev/vol/`*`disk_group`* directory.
- A volume's character device interface is located in the `/dev/rvol/`*`disk_group`* directory.

Because these interfaces support the standard UNIX `open`, `close`, `read`, `write`, and `ioctl` calls, databases, file systems, applications, and secondary swap use an LSM volume in the same manner as a disk partition as shown in Figure 1–13.

Figure 1–13: Using LSM Volumes Like Disk Partitions

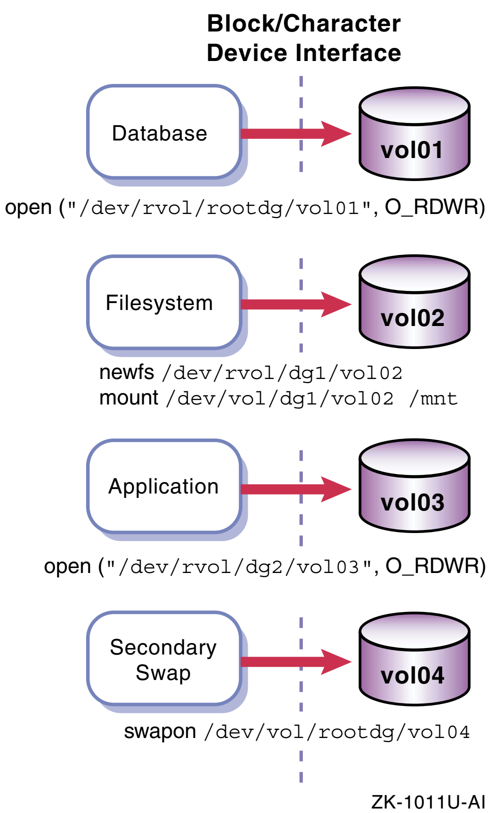

1.2 LSM Interfaces

You create, display, and manage LSM objects using any of the following interfaces:

- A Java-based graphical user interface (GUI) called LSM Storage Administrator (`lsmsa`) that displays a hierarchical view of LSM objects and their relationships.

  The Storage Administrator provides dialog boxes in which you enter information to create or manage LSM objects. Completing a dialog box can be the equivalent of entering several command-line commands. The Storage Administrator allows you to manage local or remote systems on which LSM is running. You need an LSM license to use the Storage Administrator. See Appendix A for more information on using the Storage Administrator.

- A menu-based, interactive interface called `voldiskadm`.

To perform a procedure, you choose an operation from the main menu and the `voldiskadm` interface prompts you for information. The `voldiskadm` interface provides default values when possible. You can press Return to use the default value or enter a new value or enter `?` at any time to view online help. See Appendix C and the `voldiskadm(8)` reference page for more information.

- A bit-mapped GUI called Visual Administrator (`dxlsm`) that uses the Basic X Environment.

  The Visual Administrator allows you to view and manage disks and volumes and perform limited file system administration. The Visual Administrator displays windows in which LSM objects are represented as icons. See Appendix D for more information on the Visual Administrator.

- A command-line interpreter, whereby you enter LSM commands at the system prompt. The examples in this manual use LSM commands.

In most cases, you can use the LSM interfaces interchangeably. That is, LSM objects created by one interface are manageable through and compatible with LSM objects created by other LSM interfaces. The command-line interpreter provides you with complete control and the finest granularity in creating and managing LSM objects. The other interfaces do not support all the operations available through the command line; see the relevant appendix for a description of the supported functions.

1.2.1 LSM Command-Line Interpreter

LSM provides a range of commands that allow you to display and manage LSM objects.

Table 1–1 lists the LSM commands and their functions.

Table 1–1: LSM Commands

Command	Function
`volsetup`	Initializes the LSM software
`volencap`	Sets up scripts to encapsulate disks or disk partitions into LSM volumes
`volreconfig`	Performs the encapsulation scripts set up by `volencap`, and if necessary restarts the system to complete the encapsulation
`volrootmir`	Mirrors the root and swap volumes
`volunroot`	Removes the root and swap volumes
`volmigrate`, `volunmigrate`	Migrates AdvFS domains to or from LSM volumes
`voldiskadd`	Interactively creates LSM disks

Table 1–1: LSM Commands (cont.)

Command	Function
`voldisksetup`	Adds one or more disks for use with LSM (with `-i` option)
`volassist`	Creates, mirrors, backs up, and moves volumes automatically
`voldisk`	Administers LSM disks
`voldg`	Administers disk groups
`volplex`	Administers plexes
`volume`	Administers volumes
`volsd`	Administer subdisks
`volmake`	Creates LSM objects manually
`volmirror`	Mirrors data on a disk, mirrors all unmirrored volumes in a disk group, or changes or displays the current defaults for mirroring
`voledit`	Creates, modifies, and removes LSM records
`volprint`	Displays LSM configuration information
`volsave`	Backs up the LSM configuration database
`volrestore`	Restores the LSM configuration database
`volmend`	Mends simple problems in configuration records
`volnotify`	Displays LSM configuration events
`volwatch`	Monitors LSM for failure events and performs hot-sparing if enabled
`volstat`	Displays LSM statistics
`voldctl`, `vold`, `voliod`	Controls daemon operations
`voltrace`	Traces I/O operations on volumes
`volevac`	Evacuates all volume data from a disk
`volrecover`	Synchronizes plexes after a crash or disk failure
`volinstall`	Customizes the LSM environment after LSM installation
`voldiskadm`	Starts the LSM interactive menu interface
`lsmsa`	Starts the LSM Storage Administrator GUI
`dxlsm`	Starts the LSM Visual Administrator interface

For more information on a command, see the reference page corresponding to its name. For example, for more information on the `volassist` command, enter:

```
# man volassist
```

2

Planning LSM Volumes and Disk Groups

You must plan your LSM configuration before you can use LSM volumes for applications or file systems. Planning your LSM configuration includes deciding:

- How many volumes you want, and the number and type of data plexes the volumes will use
- How many disk groups you need, and which disks you will configure in a disk group

This chapter provides information and worksheets to assist you in planning LSM volumes and disk groups. You might want to make copies of the blank worksheets for future use.

2.1 Planning LSM Volumes

Planning LSM volumes includes deciding what attributes you want the LSM volumes to have. An LSM volume has two types of attributes:

- Attributes for which you must provide a value, as described in Table 2–1.
- Attributes that are assigned a default value, which you can change, as described in Table 2–2.

Table 2–1: LSM Volume Attributes with No Default Values

Attribute	Notes
Volume name	Can be 31 alphanumeric characters but cannot include a space or slash (/). Must be unique in the disk group where you create the volume.
Volume size or length	The amount of space needed for the data in the LSM volume. You can specify volume size in sectors (the default), kilobytes, megabytes, gigabytes, or terabytes.

Table 2–2: LSM Volume Attributes with Default Values

Attribute	Notes and Default Value
Number of plexes	LSM volumes can have up to 32 plexes, with the following restrictions or recommendations: • Volumes that use a RAID 5 plex have only one data plex (RAID 5 plexes cannot be mirrored) and can have up to 31 log plexes. • Volumes that use concatenated or striped plexes can have any combination of data and log plexes for a total of 32. At least one plex should be a log plex. *Default*: One concatenated data plex, no log plex.
Log plex size	For volumes less than or equal to 1 GB that use mirror plexes, the default DRL is 65 blocks to allow for migration to a TruCluster environment. The minimum DRL size is approximately 2 blocks per GB of volume size. (You can use the minimum if you know the LSM configuration will not be used in a cluster.) For volumes that use a RAID 5 plex, the log plex size is [10 * (number of columns * data unit size)].
Plex type	A plex type is either concatenated, striped, or RAID 5. You can mirror concatenated or striped plexes. *Default*: Concatenated, no mirror. See Table 2–3 for information on choosing a plex type.
Name of the disk group where you will create the volume	A volume can be in only one disk group. *Default*: rootdg disk group.
LSM disks that the volume will use	If the volume has a striped or RAID 5 plex, each column must be of equal size and be on different disks, preferably on different buses. If the volume has mirror plexes, each data plex should use disks on different buses, and the DRL plex should be on a disk that is not used for a data plex. *Default*: LSM chooses the disks.
Usage type of the volume	Use `fsgen` for volumes that use concatenated or striped plexes and contain a file system. Use `gen` for volumes that use concatenated or striped plexes and contain data other than a file system. Use `raid5` for volumes that use a RAID 5 plex regardless of the contents of the volume. *Default*: `fsgen`.

Table 2–3 describes the benefits and trade-offs of various plex layouts and lists scenarios where one plex type might provide better performance, or be

more cost effective, than a different plex type. For optimal performance you might need to tune your system to the work load. The layout you choose depends on your specific system configuration, data availability and reliability needs, and application requirements.

Table 2–3: Choosing a Plex Type

Plex Type	Benefits and Possible Uses	Trade-offs
Concatenated	Allows you to use space on multiple disks that might otherwise be wasted. Concatenated plex can be mirrored for data redundancy. Good for volumes containing infrequently used data or data that does not change often or small volumes that can be confined to a single disk.	Possible uneven performance (hot spots, where one disk is in use by multiple applications). When mirrored, requires at least twice as much disk space (up to 32 times, depending on number of plexes).
Striped	Allows you to distribute data and therefore I/O load evenly across many disks. Good for: • Large volumes that cannot be confined to a single disk. • Applications with large read-request size. • Volumes that contain data that changes often (many writes). Striping is preferred over RAID 5 in this case, because RAID 5 imposes the overhead of calculating and writing parity data along with volume data. Striped plexes can be mirrored for data redundancy and high availability.	When mirrored, requires at least twice as much disk space (up to 32 times, depending on number of plexes).
RAID 5	Provides redundancy through parity, using fewer disks than a volume with striped mirror plexes. Provides the I/O distribution benefit of striping. Good for volumes with a high read-to-write ratio.	Depending on the I/O stripe size, performance might be slower than a volume with striped plexes due to parity calculation. The RAID 5 plex type is not supported in a cluster.

The following sections provide worksheets to assist you in planning LSM volumes depending on the type of plex you want to use. Using the information in these worksheets will help you when you create volumes as described in Chapter 4.

Note

When you create an LSM volume with the `volassist` command (the recommended and simplest method), LSM performs all the necessary calculations and creates a volume and log plexes of the appropriate sizes. The following worksheets are provided to help you approximate the space needed and ensure the disk group has enough space for the volumes you want.

2.1.1 Planning an LSM Volume That Uses a Concatenated Plex

Use the following worksheet to plan an LSM volume that uses a concatenated plex.

Figure 2–1: Worksheet for Planning a Volume with Concatenated Plexes

Attribute	Default Values	Chosen Values
Volume name	No default	
Volume size	No default	
Number of data plexes	1	
If more than one plex, DRL plex size	65 blocks for volumes less than or equal to 1 GB[a]	
Disk group name	rootdg	
Usage type	`fsgen`	
Total space required	(Volume size * number of plexes) + DRL size	

[a] For use in a cluster. Minimum DRL for standalone system is approximately 2 blocks per GB of volume size.

2.1.2 Planning an LSM Volume That Uses a Striped Plex

Use the following worksheet to plan an LSM volume that uses a striped plex.

Figure 2–2: Worksheet for Planning a Volume with Striped Plexes

Attribute	Default Values	Chosen Values
Volume name	No default	
Volume size	No default	
Data unit size	64 KB	
Number of columns	Minimum of two, based on number of disks in disk group and the volume size	
Number of data plexes	1	
If more than one plex, DRL plex size	65 blocks for volumes less than or equal to 1 GB[a]	
Disk group name	rootdg	
Usage type	`fsgen`	
Total space required	(Volume size * number of plexes) + DRL size	

[a] For use in a cluster. Minimum DRL for standalone system is approximately 2 blocks per GB of volume size.

2.1.3 Planning an LSM Volume That Uses a RAID 5 Plex

Use the following worksheet to plan an LSM volume that uses a RAID 5 plex.

Figure 2–3: Worksheet for Planning a Volume with a RAID 5 Plex

Attribute	Default Values	Chosen Values
Volume name	No default	
Volume size	No default	
Data unit size	16 KB	
Number of columns (NCOL)	Between 3 and 8 based on number if disks in disk group and the volume size	(Minimum of three)
Log plex size	10 * (data unit size * number of columns)	
Disk group name	rootdg	
Usage type	Must be `raid5`	`raid5`
Total space required	(Volume size * NCOL / (NCOL–1)) + log plex size	

2.2 Planning Disk Groups

At a minimum, you must plan the rootdg disk group, which is created when you install LSM. Planning a disk group requires that you identify:

- The space requirements for the disk group by identifying the size of volumes that you will create in the disk group, as described in Section 2.1.
- Unused storage devices to meet the space requirement of the disk group, as described in Section 2.3.

When you plan a disk group, consider the following:

- You must identify at least one unused storage device for the rootdg disk group when you install LSM.
- A disk group should have more than one storage device to ensure that there are multiple copies of the disk group's configuration database.
- To improve performance, keep the number of disks in rootdg to ten or fewer. Create additional disk groups if you:
  – Have more than ten disks to place under LSM control.
  – Might move the disk group to a different system.
- A disk group should have storage devices on different buses to improve performance and availability.
- Choose a name for the disk group carefully. There is no direct method for renaming disk groups, as there is for renaming volumes or LSM disks. To rename a disk group, you must deport it and import it with a new name. This requires stopping all activity on volumes in that disk group for the time necessary to deport and import it, which might not be acceptable or feasible in your environment.
- The following considerations apply specifically to clusters:
  – To improve performance, use the rootdg disk group only for system-related volumes such as the `cluster_root` domain and the swap devices for members. Try to keep the number of disks in rootdg to ten or fewer.
  – The disks in each disk group should be accessible by all cluster members (on a shared bus) so that all members have access to the volumes, even if one or more members are not running. If some disks are accessible only to some members, try to use those disks for data that is needed only by the member directly attached; for example, that member's swap space.

Use the worksheets in Figure 2–4 and Figure 2–5 to plan disk groups. You can make copies and fill in the information on the copies rather than in the manual. This lets you keep the disk group information with each system running LSM, for your reference. Also, because you can change your LSM configuration at any time, you can make a new copy of the blank worksheets to record your changes.

In the appropriate worksheet, enter the following:

- Under **Disk Group Information**, include any information that will help you keep track of the purpose of that disk group. For example, you might create a disk group called finance whose purpose is to contain one or more volumes that will be used by a financial application. You might create another disk group called db1, which will contain a volume used by a database.
- Under **Volume, Plex and Spare Disk Information**, include the names of all volumes in that disk group, their plex type, which disks belong to which plex and identify any spare disks that will be used to replace failed disks. See Section 3.4.4 and Section 3.4.4.1 for more information on spare disks.

  Figure 2–6 is an example of a completed worksheet.

Figure 2–4: Worksheet for Planning the rootdg Disk Group

Disk Group Information	Disks in Group	Bus/LUN Number	Disk Size	Volume, Plex, and Spare Disk Information
Name: **rootdg** Purpose:				

Figure 2–5: Worksheet for Planning Additional Disk Groups

Disk Group Information	Disks in Group	Bus/LUN Number	Disk Size	Volume, Plex, and Spare Disk Information
Name: Purpose:				

Figure 2–6 shows a consolidated example of what your disk group planning worksheets might look like when complete. Note that this example applies only to a standalone system, not a cluster.

Figure 2–6: Worksheet for Planning Disk Groups for a Standalone System (Consolidated Example)

Disk Group Information	Disks in Group	Bus/LUN Number	Disk Size	Volume, Plex, and Spare Disk Information
Name: **rootdg** Purpose: root file system and system disks.	dsk0	0	4 GB	root disk (encapsulated: rootvol plex–01)
	dsk1	0	4 GB	rootvol plex–02
	dsk4	2	4 GB	swapvol plex–01
	dsk5	2	4 GB	swapvol plex–02
	dsk16	6	4 GB	hot-spare disk
Name: **data_dg** Purpose: Database, must be redundant. Contains volume with mirrored striped plexes and DRL.	dsk6	3	18 GB	volume: db_vol plex: db_vol–01
	dsk7	3	18 GB	plex: db_vol–01
	dsk8	4	18 GB	plex: db_vol–02
	dsk9	4	18 GB	plex: db_vol–02
	dsk10	5	18 GB	plex: db_vol–03 (DRL plex)
	dsk11	5	18 GB	hot-spare disk
	dsk15	6	18 GB	hot-spare disk
Name: **finance_dg** Purpose: Financial application, must be highly available. Contains volume with RAID 5 plex (read-only application).	dsk20	7	9 GB	volume: fin_vol column: 1
	dsk25	8	9 GB	column 2
	dsk30	9	9 GB	column 3
	dsk35	10	9 GB	column 4
	dsk40	11	9 GB	column 5
	dsk45	16	9 GB	log plex
	dsk16	6	18 GB	hot-spare disk

2.3 Identifying Unused Storage Devices

Unused storage devices are unused disks, partitions, and RAID disks that LSM can initialize to become LSM disks for use in the rootdg disk group or in the other disk groups that you create.

You can also identify unused LSM disks for use in a disk group. An unused LSM disk is a storage device that you initialized for use by LSM but did not assign to a disk group.

The following sections describe how to identify unused disks, partitions, and LSM disks. See your hardware RAID documentation for information on identifying unused hardware RAID disks.

To identify unused storage devices, you can use:

- The Disk Configuration Graphical User Interface (GUI) (Section 2.3.1).
- The command-line interpreter interface (Section 2.3.2).
- The `voldisk list` command on a system where LSM is running (Section 2.3.3).

2.3.1 Using the Disk Configuration GUI to Identify Unused Disks

To identify unused disks using the Disk Configuration GUI, start the Disk Configuration interface using either of the following methods:

- From the system prompt, enter:

  ```
  # /usr/sbin/diskconfig
  ```

- From the SysMan Applications pop-up menu on the CDE Front Panel:

  1. Choose Configuration.
  2. Double click the Disk icon in the SysMan Configuration folder.

  A window titled Disk Configuration on *hostname* is displayed. This is the main window for the Disk Configuration GUI, and lists the following information for each disk:

  – The disk name, such as dsk10

  – The device model, such as RZ1CB–CA

  – The bus number for the device

  For more information about a disk, double click on the list item (or click Configure when a disk is highlighted). The Disk Configuration: Configure Partitions: window is displayed.

  This window contains:

  – A graphical representation of the disk partitions in a horizontal bar-chart format and disk information such as the disk name, the total size of the disk, and usage information.

  – A Partition Table button that you can click to display a bar chart of the current partitions in use, their sizes, and the file system in use.

– A Disk Attributes button that you can click to display values for disk attributes.

For more information about the Disk Configuration GUI, see its online help.

2.3.2 Using Operating System Commands to Identify Unused Disks

You can use the following operating system commands to identify unused disks:

1. List all the disks on the system:

   ```
   # file /dev/rdisk/dsk*c
   ```

   Information similar to the following is displayed:

```
/dev/rdisk/dsk0c:       character special (19/38) SCSI #1 "RZ1CD-CS" disk #1 (SCSI ID #0) (SCSI LUN #0)
/dev/rdisk/dsk10c:      character special (19/198) SCSI #3 "RZ1CD-CS" disk #3 (SCSI ID #5) (SCSI LUN #0)
/dev/rdisk/dsk11c:      character special (19/214) SCSI #4 "RZ1BB-CS" disk #4 (SCSI ID #0) (SCSI LUN #0)
/dev/rdisk/dsk12c:      character special (19/230) SCSI #4 "RZ1BB-CS" disk #5 (SCSI ID #1) (SCSI LUN #0)
/dev/rdisk/dsk13c:      character special (19/246) SCSI #4 "RZ1BB-CS" disk #6 (SCSI ID #2) (SCSI LUN #0)
/dev/rdisk/dsk14c:      character special (19/262) SCSI #4 "RZ1BB-CS" disk #7 (SCSI ID #3) (SCSI LUN #0)
/dev/rdisk/dsk15c:      character special (19/278) SCSI #4 "RZ1CD-CS" disk #8 (SCSI ID #4) (SCSI LUN #0)
/dev/rdisk/dsk16c:      character special (19/294) SCSI #4 "BD009635C3" disk #9 (SCSI ID #5) (SCSI LUN #0)
/dev/rdisk/dsk17c:      character special (19/310) SCSI #4 "BD009635C3" disk #10 (SCSI ID #6) (SCSI LUN #0)
/dev/rdisk/dsk18c:      character special (19/326) SCSI #5 "RZ1CD-CS" disk #11 (SCSI ID #0) (SCSI LUN #0)
/dev/rdisk/dsk19c:      character special (19/342) SCSI #5 "RZ1BB-CS" disk #12 (SCSI ID #1) (SCSI LUN #0)
/dev/rdisk/dsk1c:       character special (19/54) SCSI #1 "RZ1BB-CA" disk #2 (SCSI ID #2) (SCSI LUN #0)
/dev/rdisk/dsk20c:      character special (19/358) SCSI #5 "RZ1CB-CA" disk #13 (SCSI ID #2) (SCSI LUN #0)
/dev/rdisk/dsk21c:      character special (19/374) SCSI #5 "RZ1CB-CA" disk #14 (SCSI ID #3) (SCSI LUN #0)
/dev/rdisk/dsk22c:      character special (19/390) SCSI #5 "RZ1CF-CF" disk #15 (SCSI ID #4) (SCSI LUN #0)
/dev/rdisk/dsk23c:      character special (19/406) SCSI #5 "RZ1CF-CF" disk #8 (SCSI ID #5) (SCSI LUN #0)
/dev/rdisk/dsk24c:      character special (19/422) SCSI #5 "BD009635C3" disk #9 (SCSI ID #6) (SCSI LUN #0)
/dev/rdisk/dsk25c:      character special (19/438) SCSI #6 "RZ1BB-CS" disk #10 (SCSI ID #1) (SCSI LUN #0)
/dev/rdisk/dsk26c:      character special (19/454) SCSI #6 "RZ1CD-CS" disk #11 (SCSI ID #3) (SCSI LUN #0)
/dev/rdisk/dsk27c:      character special (19/470) SCSI #6 "RZ1CD-CS" disk #12 (SCSI ID #5) (SCSI LUN #0)
/dev/rdisk/dsk2c:       character special (19/70) SCSI #1 "RZ1CD-CS" disk #3 (SCSI ID #4) (SCSI LUN #0)
/dev/rdisk/dsk3c:       character special (19/86) SCSI #1 "RZ1CD-CS" disk #4 (SCSI ID #6) (SCSI LUN #0)
/dev/rdisk/dsk4c:       character special (19/102) SCSI #2 "RZ1BB-CS" disk #5 (SCSI ID #0) (SCSI LUN #0)
/dev/rdisk/dsk5c:       character special (19/118) SCSI #2 "RZ1CD-CS" disk #6 (SCSI ID #2) (SCSI LUN #0)
/dev/rdisk/dsk6c:       character special (19/134) SCSI #2 "RZ1CD-CS" disk #7 (SCSI ID #4) (SCSI LUN #0)
/dev/rdisk/dsk7c:       character special (19/150) SCSI #2 "RZ1CD-CS" disk #0 (SCSI ID #6) (SCSI LUN #0)
/dev/rdisk/dsk8c:       character special (19/166) SCSI #3 "RZ1BB-CA" disk #1 (SCSI ID #1) (SCSI LUN #0)
/dev/rdisk/dsk9c:       character special (19/182) SCSI #3 "RZ1CD-CS" disk #2 (SCSI ID #3) (SCSI LUN #0)
```

2. To verify if a disk or partition is unused, choose a disk from the output of the `file /dev/rdisk/dsk*c` command and enter the `disklabel` command with the name of the disk; for example:

   ```
   # disklabel dsk20c
   ```

   Disk partition information similar to the following is displayed:

   ```
   type: SCSI
   disk: RZ1CB-CA
   label:
   flags: dynamic_geometry
   bytes/sector: 512
   sectors/track: 113
   tracks/cylinder: 20
   sectors/cylinder: 2260
   cylinders: 3708
   sectors/unit: 8380080
   rpm: 7200
   interleave: 1
   ```

```
trackskew: 9
cylinderskew: 9
headswitch: 0          # milliseconds
track-to-track seek: 0  # milliseconds
drivedata: 0

8 partitions:
#          size       offset    fstype  fsize  bsize   cpg  # ~Cyl values
  a:      131072           0    unused      0      0        #      0 - 57*
  b:      262144      131072    unused      0      0        #     57*- 173*
  c:     8380080           0    unused      0      0        #      0 - 3707
  d:           0           0    unused      0      0        #      0 - 0
  e:           0           0    unused      0      0        #      0 - 0
  f:           0           0    unused      0      0        #      0 - 0
  g:     3993432      393216    unused      0      0        #    173*- 1940*
  h:     3993432     4386648    unused      0      0        #   1940*- 3707
```

See the `disklabel`(8) reference page for more information on the `disklabel` command.

3. If you are using AdvFS, display the disks in use by all domains:

```
# ls /etc/fdmns/*/*
/etc/fdmns/cluster_root/dsk7b    /etc/fdmns/root2_domain/dsk11a
/etc/fdmns/cluster_usr/dsk7g     /etc/fdmns/root_domain/dsk1a
/etc/fdmns/cluster_var/dsk7h     /etc/fdmns/usr_domain/dsk1g
/etc/fdmns/root1_domain/dsk10a
```

4. If you are using UFS, display all mounted file sets:

```
# mount
```

2.3.3 Using the LSM voldisk Command to Identify Unused Disks

When LSM starts, it obtains a list of disk device addresses from the operating system software and checks the disk labels to determine which devices are initialized for LSM use and which are not.

If LSM is running on the system, you can use the `voldisk` command to display a list of all known disks and to display detail information about a particular disk:

1. To view a list of disks, enter:

```
# voldisk list
```

Information similar to the following is displayed.

```
DEVICE       TYPE      DISK         GROUP        STATUS
dsk0         sliced    -            -            unknown
dsk1         sliced    -            -            unknown
dsk2         sliced    dsk2         rootdg       online
dsk3         sliced    dsk3         rootdg       online
dsk4         sliced    dsk4         rootdg       online
dsk5         sliced    dsk5         rootdg       online
dsk6         sliced    dsk6         dg1          online
dsk7         sliced    -            -            online
dsk8         sliced    dsk8         dg1          online
dsk9         sliced    -            -            online
```

```
dsk10        sliced    -            -            online
dsk11        sliced    -            -            online
dsk12        sliced    -            -            online
dsk13        sliced    -            -            unknown
dsk14        sliced    -            -            unknown
```

The following list describes the information in the output:

`DEVICE`
: Specifies the disk access name assigned by the operating system.

`TYPE`
: Specifies the LSM disk type (sliced, simple, or nopriv).

`DISK`
: Specifies the LSM disk media name. A dash (–) means the device is not assigned to a disk group and therefore does not have an LSM disk media name.

`GROUP`
: Specifies the disk group to which the device belongs. A dash (–) means the device is not assigned to a disk group.

`STATUS`
: An unused storage device is one that does not have a DISK name or GROUP name and has a status of unknown.

  An unused LSM disk is one that has a DISK name but has no GROUP name and a status of online or offline.

2. To display detail information about an LSM disk, enter:

   \# **voldisk list *disk***

   The following example displays information for an LSM disk called dsk5:

```
Device:    dsk5
devicetag: dsk5
type:      sliced
hostid:    servername
disk:      name=dsk5 id=942260116.1188.servername
group:     name=dg1 id=951155418.1233.servername
flags:     online ready autoimport imported
pubpaths:  block=/dev/disk/dsk5g char=/dev/rdisk/dsk5g
privpaths: block=/dev/disk/dsk5h char=/dev/rdisk/dsk5h
version:   n.n
iosize:    min=512 (bytes) max=2048 (blocks)
public:    slice=6 offset=16 len=2046748
private:   slice=7 offset=0 len=4096
update:    time=952956192 seqno=0.11
headers:   0 248
configs:   count=1 len=2993
logs:      count=1 len=453
Defined regions:
 config   priv     17-   247[   231]: copy=01 offset=000000 enabled
 config   priv    249-  3010[  2762]: copy=01 offset=000231 enabled
 log      priv   3011-  3463[   453]: copy=01 offset=000000 enabled
```

   The size of an LSM disk is displayed in blocks as the `len=` value in the `public:` row. 2048 blocks equal 1 MB.

See the `voldisk`(8) reference page for more information on the `voldisk` command.

3

Installing, Upgrading, or Uninstalling the LSM Software

This chapter describes how to:

- Prepare an existing LSM configuration on a system running Tru64 UNIX Version 4.0 for reuse when the system is upgraded to Tru64 UNIX Version 5.0 or higher (Section 3.1), which includes:
  - Increasing the size of BCLs
  - Deporting disk groups that you do not want to upgrade
  - Backing up the current LSM configuration
- Install or upgrade the LSM software (Section 3.2)
- Install the LSM license (Section 3.3), which is necessary to:
  - Create volumes with striped or RAID 5 plexes
  - Create volumes with mirror plexes
  - Use the other LSM interfaces
- Perform the following optional LSM postinstallation tasks:
  - Initialize the LSM software (Section 3.4.1)
  - Optimize the configuration database (Section 3.4.2)
  - Create an alternate boot disk (Section 3.4.3)
  - Configure the LSM hot-sparing feature (Section 3.4.4)
- Uninstall the LSM software from a system running Tru64 UNIX Version 5.0 or higher (Section 3.6)

3.1 Preparing to Upgrade LSM

If you are currently using LSM on a system running Tru64 UNIX Version 4.0 and you want to preserve your current LSM configuration for use with Tru64 UNIX Version 5.0 or higher, you must:

- Increase the size of any block-change logs (BCLs) to at least two blocks
- Optionally, deport any disk groups that you do not want to upgrade
- Back up the current LSM configuration

3.1.1 Increasing the Size of BCLs

The dirty-region logging (DRL) feature is a replacement for the block-change logging (BCL) feature that was supported in Tru64 UNIX Version 4.0. This section applies only if you are upgrading a system with an existing LSM configuration from Tru64 UNIX Version 4.0 to Version 5.0 or higher.

When you perform an upgrade installation, BCLs are automatically converted to DRLs if the BCL subdisk is at least two blocks. If the BCL subdisk is one block, logging is disabled after the upgrade installation.

Note

The conversion of BCLs to DRLs is not reversible.

Before you upgrade, increase the size of the BCLs to at least two blocks for standalone systems or 65 blocks for a TruCluster environment. If this is not possible, then after the upgrade you can enable DRL in those volumes with the `volassist addlog` command (Section 5.5.3). The `volassist addlog` command creates a DRL of 65 blocks by default.

3.1.2 Deporting Disk Groups

In Tru64 UNIX Version 5.0 and higher, LSM has an internal metadata format that is not compatible with the metadata format of LSM in Tru64 UNIX Version 4.0. If LSM detects an older metadata format during the upgrade procedure, LSM automatically upgrades the old format to the new format. If you do not want certain disk groups to be upgraded, you must deport them before you upgrade LSM.

To deport a disk group, enter:

```
# voldg deport disk_group ...
```

If you later import a deported disk group, LSM upgrades the metadata format.

3.1.3 Backing Up a Previous LSM Configuration

Backing up the LSM configuration creates a file that describes all the LSM objects in all disk groups. In case of a catastrophic failure, LSM can use this file to restore the LSM configuration.

Caution

The following procedure backs up only the configuration, not the volume data. You might also want to back up the volume data before performing the upgrade. See Section 5.4.2 for information on backing up volumes.

To back up an LSM configuration:

1. Start the backup procedure:

   ```
   # volsave [-d dir]
   ```

   Information similar to the following is displayed:

   ```
   LSM configuration being saved to /usr/var/lsm/db/LSM.date.LSM_hostidname
   LSM Configuration saved successfully to /usr/var/lsm/db/LSM.date.LSM_hostidname
   ```

   By default, LSM configuration information is saved to a time-stamped file called a description set in the `/usr/var/lsm/db` directory. In the previous example, *date* is the current date and *LSM_hostidname* is, by default, the host name. Make a note of the location and name of the file. You will need this information to restore the LSM configuration after you upgrade the LSM software and the Tru64 UNIX operating system software.

2. Optionally, confirm that the LSM configuration was saved:

   ```
   # ls /usr/var/lsm/db/LSM.date.LSM_hostidname
   ```

   Information similar to the following is displayed:

   ```
   header        rootdg.d      volboot       voldisk.list
   ```

3. Save the LSM configuration to tape or other removable media.

3.2 LSM Software Subsets

The LSM software resides in three optional subsets. These are located on the CD–ROM containing the base operating system software for the Tru64 UNIX product kit. In the following list of subset names, *nnn* indicates the operating system version:

- `OSFLSMBIN`*nnn*

  Provides the kernel modules to build the kernel with LSM drivers. This software subset supports uniprocessor, SMP, and real-time configurations. This subset requires Standard Kernel Modules.

- `OSFLSMBASE`*nnn*

  Contains the LSM administrative commands and tools required to manage LSM. This subset is mandatory if you install LSM during a Tru64 UNIX Full Installation. This subset requires LSM Kernel Build Modules.

- `OSFLSMX11`*nnn*

Contains the LSM Motif-based graphical user interface (GUI) management tool and related utilities. This subset requires the Basic X Environment.

You can install the LSM subsets either at the same time or after you install the mandatory operating system software.

If you install the system's root file system and `/usr`, `/var`, and `swap` partitions directly into LSM volumes, the LSM subsets are installed automatically.

See the *Installation Guide* for more information on installing and upgrading the LSM software and the operating system software.

Note

If a file system was configured in an LSM volume, you must start LSM and its volumes after booting the system to single-user mode, before proceeding with the Tru64 UNIX upgrade installation.

3.3 Installing the LSM License

The LSM license that comes with the base operating system allows you to create LSM volumes that use a single concatenated plex (simple volumes). All other LSM features, such as creating LSM volumes that use striped, mirrored, and RAID 5 plexes and using the LSM GUIs, require an LSM license.

The LSM license is supplied in the form of a product authorization key (PAK) called `LSM-OA`. You load the `LSM-OA` PAK into the Tru64 UNIX License Management Facility (LMF).

If you need to order an LSM license, contact your service representative. See the `lmf(8)` reference page for more information on the License Management Facility.

3.4 Performing Postinstallation Tasks

After you install or upgrade LSM:

- If applicable, initialize LSM. See Section 3.4.1 for more information.
- If you upgraded from Tru64 UNIX Version 4.0 to Tru64 UNIX Version 5.0 or higher, you might want to modify the configuration database layout to take advantage of the automatic configuration database management. See Section 3.4.2 for more information.
- Optionally, create an alternate boot disk. This involves converting the system partitions to LSM volumes and then creating a mirror plex for the

volumes. For a standalone system, see Section 3.4.3. For a TruCluster environment, see *Cluster Administration*.

- Optionally, enable the LSM hot-sparing feature. This directs LSM to transfer data (in a mirrored or RAID 5 plex) from a failed disk to a spare disk or to free disk space. See Section 3.4.4 for more information.
- If you deported any disk groups from a system running Version 4.0 before you upgraded LSM to Version 5.0A or higher, and you want to convert the disk groups after the upgrade is complete, see Section 3.4.5.

3.4.1 Initializing LSM

Initializing LSM:

- Creates the rootdg disk group.

  You should configure at least two unused disks or partitions in the rootdg disk group to ensure there are multiple copies of the LSM configuration database. You do not have to use the rootdg disk group for your volumes, but rootdg must exist before you can create other disk groups. See Chapter 2 if you need help choosing disks or partitions for the rootdg disk group.
- Reestablishes any existing LSM configuration.
- Adds entries to the `/etc/inittab` file to automatically start LSM when the system boots.
- Creates the `/etc/vol/volboot` file, which contains the host ID.
- Creates LSM files and directories. (See Section 3.5 for a description of these files and directories.)
- Starts the `vold` and `voliod` daemons.

If you performed a full installation where the root file system and `/usr`, `/var`, and `swap` partitions were installed directly into LSM volumes, or if you performed an upgrade installation on a system that was previously running LSM, then LSM is automatically initialized.

If you were running LSM previously and performed a full installation but did not install the root file system and `/usr`, `/var`, and `swap` partitions directly into LSM volumes, then you must initialize LSM.

To initialize LSM:

1. Verify that the LSM subsets are installed:

   ```
   # setld -i | grep LSM
   ```

   LSM subset information similar to the following should display, where *nnn* indicates the operating system revision:

```
OSFLSMBASEnnn  installed Logical Storage Manager (System Administration)
OSFLSMBINnnn   installed Logical Storage Manager Kernel Modules (Kernel
               Build Environment)
OSFLSMX11nnn   installed Logical Storage Manager GUI (System Administration)
```

If the LSM subsets do not display with a status of `installed`, use the `setld` command to install them. See the *Installation Guide* for more information on installing software subsets.

2. Verify LSM drivers are configured into the kernel:

   ```
   # devswmgr -getnum driver=LSM
   ```

   LSM driver information similar to the following is displayed:

   ```
   Device switch reservation list

                                           (*=entry in use)
               driver name            instance   major
     -------------------------------   --------   -----
                                 LSM          4      43
                                 LSM          3      42
                                 LSM          2      41*
                                 LSM          1      40*
   ```

   If LSM driver information is not displayed, you must rebuild the kernel using the `doconfig` command. See the *Installation Guide* for more information on rebuilding the kernel.

3. Initialize LSM with the `volsetup` command.

   Note

   To initialize LSM in a cluster, see the *Cluster Administration* manual.

   - To reestablish an existing configuration, enter:

     ```
     # volsetup
     ```

   - If there is no existing LSM configuration, specify at least two disks or partitions for LSM to use for the rootdg disk group:

     ```
     # volsetup {disk|partition} {disk|partition ...}
     ```

     For example, to initialize LSM and use disks called dsk4 and dsk5 to create the rootdg disk group, enter:

     ```
     # volsetup dsk4 dsk5
     ```

     If you omit a disk or partition name, the `volsetup` script prompts you for it. If the `volsetup` command displays an error message that the initialization failed, you might need to reinitialize the disk. See the Disk Configuration GUI online help for more information about reinitializing a disk.

To add more disks to the rootdg disk group, use the `voldiskadd` command. See Section 5.2.3 for information on adding disks to a disk group.

3.4.1.1 Verifying That LSM Is Initialized (Optional)

Normally, you do not need to verify that LSM was initialized. If the initialization fails, the system displays error messages indicating the problem.

If you want to verify that LSM is initialized, do one or more of the following:

- To verify that the disks were added to the rootdg disk group, enter:

  ```
  # volprint
  ```

  Information similar to the following is displayed. In this case, dsk4 and dsk5 are part of the rootdg disk group.

  ```
  Disk group: rootdg

  TY NAME    ASSOC      KSTATE  LENGTH    PLOFFS   STATE    TUTIL0  PUTIL0

  dg rootdg  rootdg     -       -         -        -        -       -
  dm dsk4    dsk4       -       1854536   -        -        -       -
  dm dsk5    dsk5       -       1854536   -        -        -       -
  ```

- To verify that the `/etc/inittab` file was modified to include LSM entries, enter:

  ```
  # grep LSM /etc/inittab
  ```

  Information similar to the following is displayed:

  ```
  lsmr:s:sysinit:/sbin/lsmbstartup -b /dev/console 2>&1 ##LSM
  lsm:23:wait:/sbin/lsmbstartup -n /dev/console 2>&1 ##LSM
  vol:23:wait:/sbin/vol-reconfig -n /dev/console 2>&1 ##LSM
  ```

- To verify that the `/etc/vol/volboot` file was created, enter:

  ```
  # /sbin/voldctl list
  ```

  Information similar to the following is displayed:

  ```
  Volboot file
  version: 3/1
  seqno:   0.4
  hostid:  test.abc.xyz.com
  entries:
  ```

- To verify that the `vold` daemon is enabled, enter:

  ```
  # voldctl mode
  ```

  Information similar to the following is displayed:

  ```
  mode: enabled
  ```

- To verify that two or more `voliod` daemons are running, enter:

```
# voliod
```

Information similar to the following is displayed:

```
2 volume I/O daemons are running
```

There should be one daemon for each CPU in the system or a minimum of two.

3.4.2 Optimizing the LSM Configuration Databases (Optional)

If you restored an LSM configuration on a system that you upgraded from Tru64 UNIX Version 4.0 to Tru64 UNIX Version 5.0 or higher, you can modify the configuration databases to allow LSM to automatically manage their number and placement.

Note

This procedure is an optimization and is not required.

On systems using LSM on Tru64 UNIX Version 4.0, you had to explicitly configure between four and eight disks per disk group to have enabled databases. In Version 5.0 and higher, by default all LSM disks are configured to contain copies of the database, and LSM automatically maintains the appropriate number of enabled copies. The distinction between an enabled and disabled copy is as follows:

- Disabled—The disk's private region is configured to contain a copy of the configuration database, but this copy might be dormant (inactive). LSM enables a copy as needed; for example, when a disk with an enabled copy is removed or fails.
- Enabled—The disk's private region is configured to contain a copy of the configuration database, and this copy is active. All LSM configuration changes are recorded in all each enabled copy of the configuration database as they occur.

You should configure the private regions on all your LSM disks to contain one copy of the configuration database unless you have a specific reason for not doing so, such as:

- The disk is old or slow.
- The disk is on a bus that is heavily used.
- The private region is too small (less than 4096 blocks) to contain a copy of the configuration database (such as disks that have been migrated from earlier releases of LSM, with smaller default private regions).
- There is some other significant reason why you determine the disk should not contain a copy.

Enabling the configuration database does not use additional space on the disk; it merely sets the number of enabled copies in the private region to 1.

To set the number of configuration database copies to 1, enter:

```
# voldisk moddb disk nconfig=1
```

Disk groups containing three or fewer disks should be configured so that each disk contains two copies of the configuration database to provide sufficient redundancy. This is especially important for systems with a small rootdg disk group and one or more larger secondary disk groups.

See Section 5.3.3 for more information on modifying the LSM configuration databases.

3.4.3 Creating an Alternate Boot Disk

You can use LSM to create an alternate boot disk. If the primary boot disk fails, the system uses the alternate boot disk to remain running and can also boot from the alternate disk.

To create an alternate boot disk, you must:

1. Use the LSM encapsulation procedure to configure each partition on the boot disk into an LSM volume that uses a concatenated plex. You must also encapsulate the swap space partition if it is not on the boot disk. Encapsulation converts each partition to an LSM volume.
2. Add a mirror plex to the volumes to create copies of the data in the boot disk partitions.

Note

To facilitate recovery of environments that use LSM, you can use the bootable tape utility. This utility enables you to build a bootable standalone system kernel on magnetic tape. The bootable tape preserves your local configuration and provides a basic set of the LSM commands you will use during restoration. Refer to the `btcreate`(8) reference page and the *System Administration* manual or the online help for the SysMan Menu `boot_tape` option.

3.4.3.1 Restrictions

The following restrictions apply when you encapsulate the system partitions:

- The system cannot be part of a TruCluster cluster.

Note

To create LSM volumes for the cluster root domain, use the `volmigrate` command. See `volmigrate(8)` and the *Cluster Administration* manual for more information.

- You must encapsulate the root file system and the primary swap space partition at the same time. They do not have to be on the same disk.
- The LSM volumes are created in the rootdg disk group and have the following names:
  - `rootvol`—Assigned to the volume created for the root file system. Do not change this name, move the rootvol volume out of the rootdg disk group, or change the assigned minor device number of 0.
  - `swapvol`—Assigned to the volume created for the swap space partition. Do not change this name, move the swapvol volume out of the rootdg disk group, or change the assigned minor device number of 1.
  - All other partitions are assigned an LSM volume name based on the original partition name; for example, `vol-dsk0g`.
- The partition tables for the boot disk (and swap disk, if the swap space partition is not on the boot disk) must have at least one unused partition for the LSM private region, which cannot be the `a` or `c` partitions. LSM requires only the partition-table entry; it does not need the disk space associated with the partition. If no free space is available, LSM uses space in the swap partition and reduces the swap space by the size of the private region (4096 blocks by default).

3.4.3.2 Encapsulating System Partitions

When you encapsulate the system partitions, each partition is converted to an LSM volume with a single concatenated plex. The steps to encapsulate the system partitions are the same whether you are using the UNIX File System (UFS) or the Advanced File System (AdvFS).

The encapsulation process changes the following files:

- If you are using UFS, the `/etc/fstab` file is changed to use LSM volumes instead of disk partitions.
- If you are using AdvFS, the directories in the `/etc/fdmns` directory for domains pertaining to the boot disk (for example, root_domain and usr_domain) are changed to use LSM volumes instead of disk partitions.
- The `/etc/sysconfigtab` file is changed to update the `swapdevice` entry to use LSM volumes and to set the `lsm_rootdev_is_volume` entry to 1.

- The `bootdef_dev` console variable is changed to reflect the alternate boot disk (volume).

In addition, LSM creates a private region and stores in it a copy of the configuration database. If the system partitions are on different disks (for example, the boot partitions on dsk0 and the swap partition on dsk1), LSM creates a private region on each disk. Normally, when you encapsulate a disk or partition, LSM creates only an LSM nopriv disk for the area being encapsulated. However, because of the need to be able to boot the system even if the rest of the LSM configuration is corrupted or missing, LSM creates these special-case private regions.

Note

The encapsulation procedure requires that you restart the system.

To encapsulate the system partitions:

1. Log in as `root`.
2. Identify the name of the boot disk:

   ```
   # sizer -r
   ```

   Information similar to the following is displayed:

   ```
   dsk0
   ```

3. Identify the name of the disk on which the swap space partition is located:

   ```
   # swapon -s
   ```

   Information similar to the following is displayed:

   ```
   Swap partition /dev/disk/dsk0b (default swap):
       Allocated space:        20864 pages (163MB)
       In-use space:             234 pages (  1%)
       Free space:             20630 pages ( 98%)

   Total swap allocation:
       Allocated space:        20864 pages (163.00MB)
       Reserved space:          7211 pages ( 34%)
       In-use space:             234 pages (  1%)
       Available space:        13653 pages ( 65%)
   ```

   In the previous example, the swap space partition is located in the `b` partition on disk dsk0.

4. If the swap space partition is not on the boot disk, encapsulate the boot disk and swap disk (or disks):

   ```
   # volencap boot_disk [swap_disk ...]
   ```

For example, if dsk0 is the name of the boot disk and the swap space partition is located in the `b` partition on dsk0, enter:

```
# volencap dsk0
```

Information similar to the following is displayed:

```
Setting up encapsulation for dsk0.
    - Creating simple disk dsk0d for config area (privlen=4096).
    - Creating nopriv disk dsk0a for rootvol.
    - Creating nopriv disk dsk0b for swapvol.
    - Creating nopriv disk dsk0g.

The following disks are queued up for encapsulation or use by LSM:
 dsk0d dsk0a dsk0b dsk0g

You must now run /sbin/volreconfig to perform actual encapsulations.
```

5. Optionally, send a warning to users alerting them of the impending system shutdown.
6. Perform the actual encapsulation, and enter **now** when prompted to shut down the system:

```
# volreconfig
```

Information similar to the following is displayed:

```
The system will need to be rebooted in order to continue with
LSM volume encapsulation of:
 dsk0d dsk0a dsk0b dsk0g

Would you like to either quit and defer encapsulation until later
or commence system shutdown now? Enter either 'quit' or time to be
used with the shutdown(8) command (e.g., quit, now, 1, 5): [quit]  now
```

The system shuts down, performs the encapsulation, and automatically restarts the system.

3.4.3.3 Mirroring the System Volumes

When you encapsulate the system partitions, each partition is converted to an LSM volume with a single plex. There is still only one copy of the boot disk data. To complete the process of creating an alternate boot disk, you must add a mirror plex to each system volume.

Preferably, the disks for the mirror plexes should be on different buses than the disks that contain the original system volumes. In addition, the disks you choose for the mirrors must meet the following requirements:

- Must not already be under LSM control.
- Must have a disk label with all partitions marked unused. See the `disklabel(8)` reference page for more information.

  LSM requires one unused partition for the private region, which cannot be the `a` or `c` partitions. LSM requires only the partition-table entry; it does not need the disk space associated with the partition. If no free space is available, LSM uses space in the swap partition and reduces the swap space by the size of the private region (4096 blocks by default).

- Must be as large as the partitions on the boot disk plus the size of the private region, which by default is 4096 blocks.

  If the swap space partition is not on the boot disk, you need an additional disk for the swap space mirror plex that meets the first two requirements and is as large as the swap space partition plus the size of the private region, which by default is 4096 blocks.

  See Section 2.3 if you need help choosing a disk to use for the mirror.

Note

The following procedure does not add a log plex (DRL) to the root and swap volumes, nor should you add a log plex manually. When the system restarts after a failure, it automatically recovers the rootvol volume by doing a complete resynchronization. Attaching a log plex degrades the rootvol write performance and provides no benefit in recovery time after a system failure.

To create mirror plexes, do one of the following:

- If the swap space partition is on the boot disk, enter:

  # **volrootmir -a *boot_mirror_disk***

  For example, to create the mirror plex on a disk called dsk3, enter:

  # **volrootmir -a dsk3**

- If the swap space partition is not on the boot disk, specify separate disks for the swap space mirror plex and the boot volume mirror:

  # **volrootmir -a swap=*swap_mirror_disk boot_mirror_disk***

  See the `volrootmir`(8) reference page for more information.

3.4.3.4 Displaying Information for System Volumes

To display information for the system volumes, enter:

volprint -ht

Information similar to the following is displayed:

```
Disk group: rootdg

DG NAME         NCONFIG      NLOG     MINORS   GROUP-ID
DM NAME         DEVICE       TYPE     PRIVLEN  PUBLEN   STATE
V  NAME         USETYPE      KSTATE   STATE    LENGTH   READPOL   PREFPLEX
PL NAME         VOLUME       KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID MODE
SD NAME         PLEX         DISK     DISKOFFS LENGTH   [COL/]OFF DEVICE   MODE

dg rootdg       default      default  0        942157566.1026.hostname
```

```
.
.
.
v  rootvol       root          ENABLED  ACTIVE   262144   ROUND     -
pl rootvol-01    rootvol       ENABLED  ACTIVE   262144   CONCAT    -        RW
sd root01-01p    rootvol-01    root01   0        16       0         dsk0a    ENA
sd root01-01     rootvol-01    root01   16       262128   16        dsk0a    ENA
pl rootvol-02    rootvol       ENABLED  ACTIVE   262144   CONCAT    -        RW
sd root02-02p    rootvol-02    root02   0        16       0         dsk3a    ENA
sd root02-02     rootvol-02    root02   16       262128   16        dsk3a    ENA

v  swapvol       swap          ENABLED  ACTIVE   333824   ROUND     -
pl swapvol-01    swapvol       ENABLED  ACTIVE   333824   CONCAT    -        RW
sd swap01-01     swapvol-01    swap01   0        333824   0         dsk0b    ENA
pl swapvol-02    swapvol       ENABLED  ACTIVE   333824   CONCAT    -        RW
sd swap02-02     swapvol-02    swap02   0        333824   0         dsk3b    ENA

v  vol-dsk0g     fsgen         ENABLED  ACTIVE   1450796  SELECT    -
pl vol-dsk0g-01  vol-dsk0g     ENABLED  ACTIVE   1450796  CONCAT    -        RW
sd dsk0g-01      vol-dsk0g-01  dsk0g-AdvFS 0     1450796  0         dsk0g    ENA
pl vol-dsk0g-02  vol-dsk0g     ENABLED  ACTIVE   1450796  CONCAT    -        RW
sd dsk3g-01      vol-dsk0g-02  dsk3g-AdvFS 0     1450796  0         dsk3g    ENA
```

The previous example shows that there are three volumes:

- rootvol
- swapvol
- vol–dsk0g (which contains the `/usr` partition)

Each volume has two plexes (listed in the rows labeled `pl`). Each plex uses a different subdisk (listed in the rows labeled `sd`). The first plex in each volume uses a subdisk on dsk0 (the original disk) and the second plex uses a subdisk on dsk3, indicating that the plexes were successfully mirrored onto dsk3.

The subdisks labeled root01–01p and root02–02p are phantom subdisks. Each is 16 sectors long. They provide write-protection for block 0, which prevents accidental destruction of the boot block and disk label.

3.4.3.5 Displaying Encapsulated AdvFS Domain Information

If the root file system is AdvFS, the encapsulation process automatically changes the domain information to reflect volume names instead of disk partitions.

To display the changed names:

1. Change to the `fdmns` directory:

   ```
   # cd /etc/fdmns
   ```

2. Display attributes of all AdvFS domains:

   ```
   # showfdmn *
   ```

   Information similar to the following is displayed that shows the volume name for each AdvFS domain:

```
            Id             Date Created  LogPgs  Version  Domain Name
3a5e0785.000b567c  Thu Jan 11 14:20:37 2001     512        4  root_domain

 Vol   512-Blks      Free  % Used  Cmode  Rblks  Wblks  Vol Name
  1L     524288    339936     35%     on    256    256  /dev/vol/rootdg/rootvol

            Id             Date Created  LogPgs  Version  Domain Name
3a5e078e.000880dd  Thu Jan 11 14:20:46 2001     512        4  usr_domain

 Vol   512-Blks      Free  % Used  Cmode  Rblks  Wblks  Vol Name
  1L    2879312   1703968     41%     on    256    256  /dev/vol/rootdg/vol-dsk0g

              Id             Date Created  LogPgs  Version  Domain Name
3a5e0790.0005b501  Thu Jan 11 14:20:48 2001     512        4  var_domain

 Vol   512-Blks      Free  % Used  Cmode  Rblks  Wblks  Vol Name
  1L    2879312   2842160      1%     on    256    256  /dev/vol/rootdg/vol-dsk0h
```

3.4.3.6 Displaying Encapsulated UFS File System Information

If the root file system is UFS, the encapsulation process automatically changes the mount information to reflect volume names instead of disk partitions.

To display the volume names for the root file system, enter:

```
# mount
```

Information similar to the following is displayed. File systems of the form `/dev/vol/`*`disk_group`*`/`*`volume`* indicate that the file system is encapsulated into LSM volumes.

```
/dev/vol/rootdg/rootvol on / type ufs (rw)
/dev/vol/rootdg/vol-dsk2g on /usr type ufs (rw)
/proc on /proc type procfs (rw)
```

3.4.3.7 Displaying Encapsulated Swap Volume Information

To display the volume information for the swap space, enter:

```
# swapon -s
```

Information similar to the following is displayed:

```
Swap partition /dev/vol/rootdg/swapvol (default swap):
    Allocated space:        20864 pages (163MB)
    In-use space:             234 pages (  1%)
    Free space:             20630 pages ( 98%)

Total swap allocation:
    Allocated space:        20864 pages (163.00MB)
    Reserved space:          7211 pages ( 34%)
    In-use space:             234 pages (  1%)
    Available space:        13653 pages ( 65%)
```

3.4.4 Automatic Data Relocation Feature (Hot-Sparing)

You can enable the LSM hot-sparing feature to configure LSM to automatically relocate data from a failed disk in a volume that uses either a RAID 5 plex or mirrored plexes. LSM relocates the data to either a reserved disk that you configured as a spare disk or to free disk space in the disk group. LSM does not use a spare disk for normal data storage unless you specify otherwise.

During the hot-sparing procedure, LSM:

- Sends mail to the `root` account (and other specified accounts) with notification about the failure and identifies the affected LSM objects.
- Determines which LSM objects to relocate.
- Relocates the LSM objects from the failed disk to a spare disk or to free disk space in the disk group. However, LSM will not relocate data if redundancy cannot be preserved. For example, LSM will not relocate data to a disk that contains a mirror of the data.
- Updates the configuration database with the relocation information.
- Ensures that the failed disk space is not recycled as free disk space.
- Sends mail to the `root` account (and other specified accounts) about the action taken.

If you choose not to use the hot-spare feature, you must investigate and resolve disk failures manually. See Section 6.5 for more information.

3.4.4.1 Enabling the Hot-Sparing Feature

The hot-sparing feature is part of the `volwatch` daemon. The `volwatch` daemon has two modes:

- Mail-only, which is the default. This setting notifies you of a problem but does not perform hot-sparing. You can reset the daemon to this mode with the `-m` option.
- Mail-and-spare, which you set with the `-s` option.

You can specify mail addresses with either option.

To enable the hot-sparing feature, enter:

```
# volwatch -s [mail-address...]
```

Note

Only one `volwatch` daemon can run on a system or cluster node at any time. The daemon's setting applies to the entire system or node; you cannot specify some disk groups to use hot-sparing but not others.

To return the volwatch daemon to mail-only mode, enter:

```
# volwatch -m [mail-address...]
```

3.4.4.2 Configuring and Deconfiguring a Spare Disk

You should configure at least one spare disk in each disk group that contains volumes with mirror plexes or a RAID 5 plex.

To configure a disk as a spare, enter:

```
# voledit [-g disk_group] set spare=on disk
```

For example, to configure a spare disk called dsk5 in the rootdg disk group, enter:

```
# voledit set spare=on dsk5
```

To deconfigure a spare disk, enter:

```
# voledit [-g disk_group] set spare=off disk
```

For example, to deconfigure a spare disk called dsk5 in the rootdg disk group, enter:

```
# voledit set spare=off dsk5
```

3.4.4.3 Setting Up Mail Notification for Exception Events

The `volwatch` daemon monitors LSM for exception events. If an exception event occurs, mail is sent to the `root` account and to other accounts that you specify:

- When you use the `rcmgr` command to set the `VOLWATCH_USERS` variable in the `/etc/rc.config.common` file. See the `rcmgr`(8) reference page for more information on the `rcmgr` command.
- On the command line with the `volwatch` command.

There is a 15-second delay before the event is analyzed and the message is sent. This delay allows a group of related events to be collected and reported in a single mail message.

Example 3–1 shows a sample mail notification sent when LSM detects an exception event.

Example 3–1: Sample Mail Notification

```
Failures have been detected by the Logical Storage Manager:

failed disks:
disk
.
.
.

failed plexes:
plex
.
.
.

failed log plexes:
plex
.
.
.

failing disks:
disk
.
.
.

failed subdisks:
subdisk
.
.
.

The Logical Storage Manager will attempt to find spare disks,
relocate failed subdisks and then recover the data in the failed plexes.
```

The following describes the sections of the mail notification:

- The *disk* under `failed disks` specifies disks that appear to have failed completely.
- The *plex* under `failed plexes` shows plexes that are detached due to I/O failures experienced while attempting to do I/O to subdisks they contain.
- The *plex* under `failed log plexes` indicates RAID 5 or dirty region log (DRL) plexes that have experienced failures.
- The *disk* under `failing disks` indicates a partial disk failure or a disk that is in the process of failing. When a disk has failed completely, the same *disk* appears under both `failed disks` and `failing disks`.
- The *subdisk* under `failed subdisks` indicates a subdisk in a RAID 5 volume that has been detached due to I/O errors.

Example 3–2 shows the mail message sent if a disk completely fails.

Example 3–2: Complete Disk Failure Mail Notification

```
To: root
Subject: Logical Storage Manager failures on servername.com

Failures have been detected by the Logical Storage Manager

failed disks:
 disk02

failed plexes:
  home-02
  src-02
  mkting-01

failing disks:
disk02
```

This message shows that a disk called disk02 was failing, then detached by a failure and that plexes called home-02, src-02 and mkting-01 were also detached (probably due to the disk failure).

Example 3–3 shows the mail message sent if a disk partially fails.

Example 3–3: Partial Disk Failure Mail Notification

```
To: root
Subject: Logical Storage Manager failures on servername.com

Failures have been detected by the Logical Storage Manager:

failed disks:
 disk02

failed plexes:
  home-02
  src-02
```

Example 3–4 shows the mail message sent if data relocation is successful and data recovery is in progress.

Example 3–4: Successful Data Relocation Mail Notification

```
Volume volume Subdisk subdisk relocated to new_subdisk,
but not yet recovered.
```

If the data recovery is successful, the following message is sent:

```
Recovery complete for volume volume in disk group disk_group.
```

If the data recovery is unsuccessful, the following message is sent:

```
Failure recovering volume in disk group disk_group.
```

Example 3–5 shows the mail message sent if relocation cannot occur because there is no spare or free disk space.

Example 3–5: No Spare or Free Disk Space Mail Notification

```
Relocation was not successful for subdisks on disk disk
in volume volume in disk group disk_group.
No replacement was made and the disk is still unusable.

The following volumes have storage on disk:

volume
.
.
.

These volumes are still usable, but the redundancy of those
volumes is reduced. Any RAID5 volumes with storage on the
failed disk may become unusable in the face of further failures.
```

Example 3–6 shows the mail message that is sent if data relocation fails.

Example 3–6: Data Relocation Failure Mail Notification

```
Relocation was not successful for subdisks on disk disk in
volume volume in disk group disk_group.
No replacement was made and the disk is still unusable.

error message
```

In this output, *error message* is a message indicating why the data relocation failed.

Example 3–7 shows the mail message sent if volumes not using RAID 5 plexes are made unusable due to disk failure.

Example 3–7: Unusable Volume Mail Notification

```
The following volumes:

volume
.
.
.

have data on disk but have no other usable mirrors
on other disks. These volumes are now unusable and the data on
them is unavailable. These volumes must have their data restored.
```

Example 3–8 shows the mail message sent if volumes using RAID 5 plexes are made unusable due to disk failure.

Example 3–8: Unusable RAID 5 Volume Mail Notification

```
The following RAID5 volumes:

volume
.
.
.

have storage on disk and have experienced
other failures. These RAID5 volumes are now unusable
and data on them is unavailable. These RAID5 volumes must
have their data restored.
```

3.4.4.4 Moving Relocated LSM Objects

When data is moved by the hot-sparing feature, the new locations of LSM objects might not provide the same performance or have the same data layout that existed before. After hot-sparing occurs, you might want to move the relocated LSM objects to improve performance, to keep the spare disk space free for future hot-sparing needs, or to restore the LSM configuration to its previous state.

Note

This procedure assumes you have identified and initialized a new disk to replace the hot-spare disk. See Section 4.1.2 for more information on adding disks for LSM use. See Section 6.5.5 for information on replacing a failed disk.

To move a subdisk that was relocated as the result of a hot-sparing procedure:

1. Note the characteristics of the LSM objects before they were relocated.

   This information is available from the mail notification sent to `root`. For example, look for a mail notification similar to the following:

   ```
   To: root
   Subject: Logical Storage Manager failures on host teal

   Attempting to relocate subdisk disk02-03 from plex home-02.
   Dev_offset 0 length 1164 dm_name disk02 da_name dsk2.
   The available plex home-01 will be used to recover the data.
   ```

2. Note the new location for the relocated LSM object.

   This information is available from the mail notification sent to `root`. For example, look for a mail notification similar to the following:

   ```
   To: root
   Subject: Attempting LSM relocation on host teal

   Volume home Subdisk disk02-03 relocated to disk05-01,
   but not yet recovered.
   ```

3. Move the relocated data to the desired location:

   ```
   # volevac [-g disk_group] spare_disk new_disk
   ```

4. Move the LSM volume from the hot-spare disk to the new disk. The `!` prefix indicates the source disk. Use the appropriate shell quoting convention to correctly interpret the `!`.

   ```
   # volassist [-g disk_group] move volume !hot_spare new_disk
   ```

3.4.5 Importing and Converting Tru64 UNIX Version 4.0 Disk Groups

After you upgrade a system, you can import and convert disk groups that you deported before the upgrade.

Note

This section applies only to disk groups deported from a system running Tru64 UNIX Version 4.0, to be imported to a system running Tru64 UNIX Version 5.0 or higher.

Importing the disk group upgrades its metadata format to the current format. Converting the disk group changes all volumes that use BCLs to use DRLs instead. Use the `vollogcnvt` utility to manually perform this conversion.

The `vollogcnvt` utility runs automatically when the system is started, converting all volumes possible in imported disk groups. (See the `vollogcnvt`(8) reference page for more information.)

You can manually import a disk group, determine whether its volumes need to have BCLs converted to DRLs, and then run the `vollogcnvt` utility without having to shut down and restart the system.

To upgrade disk groups to the current metadata format and convert volumes from BCL to DRL:

1. Import the disk group with the conversion option:

   ```
   # voldg -o convert_old import disk_group
   ```

   The disk group is imported. If any volumes in the disk group use BCLs, a message similar to the following is displayed:

   ```
   lsm:voldg:WARNING:Logging disabled on volume. Need to convert to DRL.
   lsm:voldg:WARNING:Run the vollogcnvt command to automatically convert logging.
   ```

   All the volumes in the disk group are usable, but logging is disabled for volumes that previously used BCL. If a disk in the volume fails or the system crashes, the entire volume will be resynchronized when the disk is replaced or the system restarts.

2. Convert the volumes from BCL to DRL:

   ```
   # vollogcnvt [-o disk_group]
   ```

   All possible volumes are converted from using BCL to DRL.

   If the volume cannot be converted, logging is disabled but the volume is usable, and data continues to be written to all mirrors (plexes) in the volume.

3. To restore logging, remove the disabled BCL subdisk and add a new DRL to the volume:

   a. Identify the disabled BCL subdisk:

      ```
      # volprint [-o disk_group] volume
      ```

   b. Remove the BCL subdisk:

      ```
      # volsd [-o rm] dis subdisk
      ```

See Section 5.5.3 for information on adding a DRL to a volume.

3.5 LSM Files, Directories, Device Drivers, and Daemons

After you install and initialize LSM, several new files, directories, device drivers, and daemons are present on the system. These are described in following sections.

3.5.1 LSM Files

The /dev directory contains the device special files (Table 3–1) that LSM uses to communicate with the kernel.

Table 3–1: LSM Device Special Files

Device Special File	Function
/dev/volconfig	Allows the vold daemon to make configuration requests to kernel
/dev/volevent	Used by the voliotrace command to view and collect events
/dev/volinfo	Used by the volprint command to collect LSM object status information
/dev/voliod	Provides an interface between the volume extended I/O daemon (voliod) and the kernel

3.5.2 LSM Directories

The /etc/vol directory contains the volboot file and the subdirectories (Table 3–2) for LSM use.

Table 3–2: LSM /etc/vol Subdirectories

Directory	Function
`reconfig.d`	Provides temporary storage during encapsulation of existing file systems. Instructions for the encapsulation process are created here and used during the reconfiguration.
`tempdb`	Used by the volume configuration daemon (`vold`) while creating the configuration database during startup and while updating configuration information.
`vold_diag`	Creates a socket portal for diagnostic commands to communicate with the `vold` daemon.
`vold_request`	Provides a socket portal for LSM commands to communicate with the `vold` daemon.

The `/dev` directory contains the subdirectories (Table 3–3) for volume block and character devices.

Table 3–3: LSM Block and Character Device Subdirectories

Directory	Contains
`/dev/rvol`	Character device interfaces for LSM volumes.
`/dev/vol`	Block device interfaces for LSM volumes.

3.5.3 LSM Device Drivers

There are two LSM device drivers:

- `volspec`—The volume special device driver

  Communicates with the LSM device special files. This is not a loadable driver; it must be present at boot time.

- `voldev`—The volume device driver

  Communicates with LSM volumes. Provides an interface between LSM and the physical disks.

3.5.4 LSM Daemons

There are two LSM daemons:

- `vold`—The Volume Configuration Daemon. This daemon is responsible for maintaining configurations of disks and disk groups. It also:
  - Takes requests from other utilities for configuration changes
  - Communicates change requests to the kernel
  - Modifies configuration information stored on disk

   - Initializes LSM when the system starts

- `voliod`—The Volume Extended I/O Daemon. This daemon performs the functions of a utility and a daemon. As a utility, `voliod`:
   - Returns the number of running volume I/O daemons
   - Starts more daemons when necessary
   - Removes some daemons from service when they are no longer needed

  As a daemon, `voliod`:
   - Schedules I/O requests that must be retried
   - Schedules writes that require logging (for DRL and RAID 5 log plexes)

3.6 Uninstalling the LSM Software

This section describes how to completely remove the LSM software from a system. This process involves:

- Backing up user data
- Unencapsulating disks or data
- Removing LSM objects and the software subsets

Caution

Uninstalling LSM causes any current data to be lost. You should unencapsulate and back up any needed data before proceeding.

---

1. Reconfigure any system file systems and swap space so they are no longer on an LSM volume.
   a. If root and swap are configured under LSM, enter the `volunroot` command and restart the system.

      If additional swap space was configured using LSM volumes, remove those volumes (Section 5.4.6).
   b. Unencapsulate the `/usr` and `/var` file systems if these are configured under LSM. See Section 5.4.6.2 if `/usr` and `/var` are encapsulated under LSM.
2. Unmount any other file systems that are using LSM volumes so all LSM volumes can be closed.
   a. Update the `/etc/fstab` file if necessary so that it no longer mounts any file systems on an LSM volume.
   b. Stop applications that are using raw LSM volumes and reconfigure them so that they no longer use LSM volumes.

3. Identify the disks that are currently configured under LSM:

```
# voldisk list
```

4. Restart LSM in disabled mode:

```
# vold -k -r reset -d
```

This command fails if any volumes are open.

5. Stop all LSM volume and I/O daemons:

```
# voliod -f set 0
# voldctl stop
```

6. Update the disk labels for the disks under LSM. See step 3.

   - For each LSM sliced disk, repartition and update the disk labels:

   ```
   # disklabel -rw disk
   ```

   - For each LSM simple disk, update the partition's `fstype` field to unused:

   ```
   #  disklabel -s dsknc unused
   ```

   - For each LSM nopriv disk, update the partition's `fstype` field to either unused or the appropriate value, depending on whether the partition still contains valid data.

     For example, if dsk2g was an LSM nopriv disk that still contains a valid UNIX File System and dsk2h was an LSM nopriv disk that no longer contains valid data, enter:

   ```
   # disklabel -s dsk2g 4.2BSD
   # disklabel -s dsk2h unused
   ```

7. Remove the LSM directories:

```
# rm -r /etc/vol /dev/vol /dev/rvol /etc/vol/volboot
```

8. Delete the following LSM entries in the `/etc/inittab` file:

```
lsmr:s:sysinit:/sbin/lsmbstartup -b </dev/console >/dev/console 2>&1 ##LSM
lsm:23:wait:/sbin/lsmbstartup </dev/console >/dev/console 2>&1 ##LSM
vol:23:wait:/sbin/vol-reconfig -n </dev/console >/dev/console 2>&1 ##LSM
```

9. Display the installed LSM subsets:

```
# setld -i | grep LSM
```

10. Delete the installed LSM subsets:

```
# setld -d OSFLSMBASEnnn OSFLSMBINnnn OSFLSMCLSMTOOLSnnn
```

11. To deconfigure LSM from the kernel, replace the `pseudo-device lsm 1` entry in the `/sys/conf/`*hostname* file to `pseudo-device lsm 0`.

    You can make this change either prior to running the `doconfig` command or while running the `doconfig` command; for example:

```
# doconfig -c hostname
```

12. Copy the new kernel to `root` (/) and restart the system by entering the following commands:

```
# cp /sys/RIO/vmunix /
# shutdown now
```

When the system restarts, LSM will no longer be installed.

4 Creating LSM Disk Groups and Volumes

An LSM volume is an object that represents a hierarchy of LSM objects that allocates space to, and stores data for, a file system or application. You create an LSM volume differently depending on whether the volume is for a new file system or application or an existing file system or application.

This chapter describes how to:

- Create a disk group and check a disk group for space (Section 4.1)
- Create an LSM volume for new data (Section 4.2)
- Configure UFS or AdvFS file systems to use an LSM volume (Section 4.3)
- Create an LSM volume for existing data (Section 4.4)

Use the information from the worksheets you filled out in Chapter 2 to create disk groups and LSM volumes.

4.1 Creating Disk Groups

You must create an LSM volume in a disk group. By default, LSM creates volumes in the rootdg disk group, which was created when you installed LSM. You can create all LSM volumes in the rootdg disk group or you can create other disk groups. The following sections describe how to:

- Display disk group information
- Create a disk group
- Add disks to a disk group
- Create a backup copy of the disk label information

4.1.1 Displaying Disk Group Information

To display information about the rootdg disk group and other disk groups, enter:

```
# voldisk list
```

Information similar to the following is displayed:

```
DEVICE       TYPE      DISK        GROUP       STATUS
dsk0         sliced    -           -           unknown
dsk1         sliced    -           -           unknown
dsk2         sliced    dsk2        rootdg      online
dsk3         sliced    dsk3        rootdg      online
dsk4         sliced    dsk4        rootdg      online
dsk5         sliced    dsk5        rootdg      online
dsk6         sliced    dsk6        dg1         online
dsk7         sliced    dsk7        dg1         online
dsk8         sliced    -           -           unknown
```

where:

`DEVICE` — Specifies the disk access name assigned by the operating system software.

`TYPE` — Specifies the LSM disk type: sliced, simple, or nopriv.

`DISK` — Specifies the LSM disk media name. An LSM disk media name displays only if the disk is in a disk group.

`GROUP` — Specifies the disk group to which the device belongs. A group name displays only if a disk is in a disk group.

`STATUS` — Specifies the status of the LSM device. The status is one of the following:

- `online`—The device was initialized for LSM use and is in use.
- `offline`—The device was initialized for LSM use, but is not available.
- `unknown`—LSM detected the device, but it was not initialized for LSM use (has no disk media name and is not part of a disk group).
- `error` — The disk is detected but has experienced I/O errors.
- `failed was` — An LSM disk media name exists but the disk is not associated with a `DEVICE`. Displays the last device associated with this name.

To display the total usable space in a disk group, enter:

volassist [**-g** ***disk_group***] **maxsize**

The following command line displays the available space in a disk group called dg1:

volassist -g dg1 maxsize

Information similar to the following is displayed:

```
Maximum volume size: 6139904 (2998Mb)
```

4.1.2 Creating a Disk Group or Adding Disks to a Disk Group

The `voldiskadd` script is an interactive script that lets you:

- Initialize disks or disk partitions for exclusive use by LSM
- Create a disk group
- Add disks to a disk group

Note

By default, LSM initializes each disk with one copy of the configuration database. If a disk group will have fewer than four disks, you should initialize each disk to have two copies of the disk group's configuration database to ensure that the disk group has multiple copies in case one or more disks fail. You must use the `voldisksetup` command to initialize disks with more than one copy of the configuration database. See Section 5.1.1 for more information.

If you specify an uninitialized disk, LSM initializes the disk as an LSM sliced disk. If you specify a partition name, LSM initializes the partition as an LSM simple disk. You can specify several disks and disk partitions at once, separated by a space; for example:

```
# voldiskadd dsk3 dsk4a dsk5 dsk6g
```

After you initialize a disk or disk partition, LSM writes a new disk label and the disk or disk partition becomes an LSM disk for exclusive use by LSM.

The `voldiskadd` script prompts you for the following information:

- A disk group name

  If you are creating a disk group, the disk group name must be unique and can contain up to 31 alphanumeric characters that cannot include spaces or the forward slash (/).

- A disk media name for each disk you configure in the disk group

  You can use the default disk media name or you can assign a disk media name of up to 31 alphanumeric characters that cannot include spaces or the forward slash (/).

- Whether the disk should be a spare disk for the disk group

  A spare disk is a disk initialized by LSM, but used only as a replacement disk if a disk that contains a mirror or RAID 5 plex fails. See Section 3.4.4 for more information about how LSM uses spare disks. For the best protection, configure at least one spare disk in each disk group that contains mirror or RAID 5 plexes.

The following example uses a disk called dsk9 to create a disk group called dg1:

```
# voldiskadd dsk9
```

Information similar to the following is displayed:

```
Add or initialize disks

Menu: VolumeManager/Disk/AddDisks

  Here is the disk selected.

  dsk9

Continue operation? [y,n,q,?] (default: y) y

  You can choose to add this disk to an existing disk group, a
  new disk group, or leave the disk available for use by future
  add or replacement operations. To create a new disk group,
  select a disk group name that does not yet exist. To leave
  the disk available for future use, specify a disk group name
  of "none".

Which disk group [<group>,none,list,q,?] (default: rootdg) dg1

  There is no active disk group named dg1.

Create a new group named dg1? [y,n,q,?] (default: y) y

  The default disk name that will be assigned is:

  dg101

Use this default disk name for the disk? [y,n,q,?] (default: y) y

Add disk as a spare disk for dg1? [y,n,q,?] (default: n) n

  A new disk group will be created named dg1 and the selected disks
  will be added to the disk group with default disk names.
  dsk9

Continue with operation? [y,n,q,?] (default: y) y

  The following disk device has a valid disk label, but does
  not appear to have been initialized for the Logical Storage
  Manager. If there is data on the disk that should NOT be
  destroyed you should encapsulate the existing disk partitions
  as volumes instead of adding the disk as a new disk.

  dsk9

Initialize this device? [y,n,q,?] (default: y) y

  Initializing device dsk9.

  Creating a new disk group named dg1 containing the disk
  device dsk9 with the name dg101.

Goodbye.
```

4.1.3 Creating a Backup Copy of the Disk Label Information

It is highly recommended that you create a backup copy of the updated disk label information for each LSM disk.

Having this information will simplify the process of replacing a failed disk, by allowing you to copy the failed disk's attributes to a new disk. Once a disk fails, you cannot read its disk label, and you cannot copy that information to a new disk.

To create a file that contains a backup copy of the disk label information, enter:

```
# disklabel dskn > file
```

See the `disklabel`(8) reference page for more information on the `disklabel` command.

4.2 Creating an LSM Volume for New Data

To create an LSM volume for a new file system or application, use the `volassist` command. The `volassist` command finds the necessary space within the disk group and creates all the objects for the volume. You must specify a volume name and length (size) on the command line.

You can specify values for other LSM volume attributes on the command line or in a text file that you create. If you do not specify a value for an attribute, LSM uses a default value.

To display the default values for volume attributes, enter:

```
# volassist help showattrs
```

Information similar to the following is displayed:

```
#Attributes:
 layout=nomirror,nostripe,span,nocontig,raid5log,noregionlog,diskalign,nostorage
 mirrors=2 columns=0 nlogs=1 regionlogs=1 raid5logs=1
 min_columns=2 max_columns=8
 regionloglen=0 raid5loglen=0 logtype=region
 stripe_stripeunitsize=64 raid5_stripeunitsize=16
 usetype=fsgen diskgroup= comment="" fstype=
 user=0 group=0 mode=0600
 probe_granularity=2048
 alloc=
 wantalloc=
 mirror=
```

Some volume attributes have several options to define them. Some options define an attribute globally, while others define an attribute for a specific plex type. For example, you can specify the size of a stripe data unit using the `stripeunit` option for both striped or RAID 5 plexes, the `stripe_stripeunitsize` option for striped plexes, or the `raid5_stripeunitsize` option for RAID 5 plexes.

See the volassist(8) reference page for a complete list of attributes. Table 4–1 describes some of the common attributes for which you can specify a value.

Table 4–1: Common LSM Volume Attributes

Attribute Description	Attribute Options
Plex type	layout={concatenated|striped|raid5}
Usage type	-U {fsgen|raid5|gen}
Whether or not to create mirrors, and if so how many	mirror={number|yes|no}
Whether or not to use a Dirty Region Log (DRL) plex for mirrored plexes	logtype={drl|none}
Size of the stripe width for a striped or RAID 5 plex	stripeunit=*data_unit_size*
Number of columns for a striped or RAID 5 plex	nstripe=*number_of_columns*

Creating a text file that specifies many of these attributes is useful if you create many LSM volumes that use the same nondefault values for attributes. Any attribute that you can specify on the command line can be specified on a separate line in the text file. By default, LSM looks for the /etc/default/volassist file when you create an LSM volume. If you created an /etc/default/volassist file, LSM creates each volume using the attributes that you specify on the command line and in the /etc/default/volassist file.

Example 4–1 shows a text file called /etc/default/volassist that creates an LSM volume using a four-column striped plex with two mirrors, a stripe width of 32 KB, and no log.

Example 4–1: LSM Volume Attribute Defaults File

```
# LSM Vn.n
# volassist defaults file. Use '#' for comments
# number of stripes
nstripe=4
# layout
layout=striped
# mirroring
nmirror=2
# logging
logtype=none
# stripe size
stripeunit=32k
```

For example, to create an LSM volume using the attributes in the /etc/default/volassist file, enter:

```
# volassist make volume length
```

To specify a file other than the /etc/default/volassist file, you must use the volassist command with the -d option followed by the name of the file. If you use the -d option, LSM creates the volume using the attributes that you specify on the command line and in the named file.

For example, to create an LSM volume using the attributes in a file other than the /etc/default/volassist file, enter:

```
# volassist make volume length -d filename
```

The following lists the order in which LSM assigns values to attributes:

1. Values on the command line
2. Values in a file that you specify by using the volassist -d option
3. Values in the /etc/default/volassist file
4. Default values

4.2.1 Creating an LSM Volume Using a Concatenated Plex

Creating an LSM volume that uses a concatenated plex can be a multiple-step process. Step 1 is required, and the others are required only if you want to mirror the plex. To increase performance for mirror plexes, you can specify the disks for the data plexes and the DRL plex to ensure that LSM creates these plexes on different disks, preferably on different buses.

To create an LSM volume that uses a concatenated plex:

1. Create a volume with a single plex, optionally specifying disks:

   ```
   # volassist [-g disk_group] make volume length [disks]
   ```

   The following example creates a 3 GB volume called vol2 that uses disks dsk2, dsk3, and dsk4 in a disk group called dg1:

   ```
   # volassist [-g dg1] make vol2 3g dsk2 dsk3 dsk4
   ```

   The volume is created and started. If you want to mirror the plex, continue with step 2.

2. Add a mirror plex to the volume, specifying disks not used in the first data plex and preferably on different buses.

   You can use the init=active option to prevent LSM from synchronizing the plexes. Plex synchronization is not necessary because the volume does not yet contain any data. Use this option only if the application that will use the volume always writes a block before it reads from that block.

```
# volassist [-g disk_group] mirror volume [init=active] \
layout=nolog disks
```

The following example creates a mirror plex for the same volume, using disks dsk5, dsk6, and dsk7:

```
# volassist -g dg1 mirror vol2 init=active \
layout=nolog dsk5 dsk6 dsk7
```

Note

Because two plexes are used in the volume, 6 GB of free space is needed. Each plex uses 3 GB of disk space.

3. Add a DRL plex to the volume, specifying a disk that is not used by one of the data plexes:

```
#  volassist addlog volume disk
```

The following example adds a DRL plex to vol2 on a disk called dsk10:

```
#  volassist addlog vol2 dsk10
```

The volume is ready for use.

4.2.2 Creating an LSM Volume Using a Striped Plex

Creating an LSM volume that uses a striped plex can be a multiple-step process, depending on whether you want mirrors and a DRL. Step 1 creates a volume with one plex, and the next steps add a mirror and a DRL to the volume. To increase performance for the volume, you can specify the disks for each plex to ensure that LSM creates these plexes on different disks, preferably on different buses.

Note

In general, you should not use LSM to stripe data if you also use a hardware controller to stripe data. In some specific cases such a configuration can improve performance but only if:

- Most of the volume I/O requests are large (≥ 1 MB).
- The LSM volume is striped over multiple RAID sets on different controllers.
- The LSM stripe size is a multiple of the full hardware RAID stripe size.

The number of LSM columns in each plex in the volume should be equal to the number of hardware RAID controllers. See your hardware RAID documentation for information about how to choose the best number of columns for the hardware RAID set.

By default, the `volassist` command creates columns for a striped plex on disks in alphanumeric order, regardless of their order on the command line.

To improve performance, you might want to create columns using disks on different buses. See Section 4.2.3 for more information about specifying the disk order for columns in a striped plex.

To create an LSM volume that uses striped plexes:

1. Create a volume with a single plex, optionally specifying disks, preferably on different buses:

   ```
   # volassist [-g disk_group] make volume length \
   layout=stripe [nstripe=number_of_columns] \
   [stripeunit=data_unit_size] [disks]
   ```

   The following example creates a 4 GB volume called vol_stripe that uses disks dsk2, dsk3 and dsk4 to create a three-column striped plex in a disk group called dg1:

   ```
   # volassist -g dg1 make vol_stripe 4g \
   layout=stripe nstripe=3 dsk2 dsk3 dsk4
   ```

   The volume is created and started. If you want to add a mirror plex, continue with step 2.

2. Add a mirror plex to the volume, specifying disks not used in the first data plex and preferably on a different bus.

   You can use the `init=active` option to prevent LSM from synchronizing the plexes. Plex synchronization is not necessary because the volume does not yet contain any data. Use this option only if the application that will use the volume always writes a block before it reads from that block.

   ```
   # volassist [-g disk_group] mirror volume \
   [init=active] layout=nolog disks
   ```

   The following example creates a mirror plex for the volume vol_stripe, using disks dsk5, dsk6, and dsk7:

   ```
   # volassist -g dg1 mirror vol_stripe \
   init=active layout=nolog dsk5 dsk6 dsk7
   ```

Note

Because two plexes are used in the volume, 8 GB of free space is needed. Each plex uses 4 GB of disk space.

3. Add a DRL plex to the volume, specifying a disk that is not used by one of the data plexes:

   ```
   # volassist addlog volume disk
   ```

   The following example adds a DRL plex to vol_stripe on a disk called dsk10:

   ```
   # volassist addlog vol_stripe dsk10
   ```

   The volume is ready for use.

4.2.3 Creating an LSM Volume Using a Striped Plex (on Different Buses)

By default, LSM creates columns for a striped plex on the first available disks it finds in the disk group. This might result in a volume with columns that use disks on the same bus.

You can improve performance by creating a striped plex with columns that use disks on different buses. To do so, you must create the subdisks for each column.

Note

Each column of subdisks should be the same size and be a multiple of the data unit size so there is no wasted space. For example, a data unit size (stripe width) of 64 KB for a striped plex corresponds to 128 blocks (sectors), so the total of the subdisks in each column should be a multiple of 128.

If each column comprises one subdisk (the typical configuration), then the subdisk size should be a multiple of 128. If a column comprises two subdisks, one subdisk can be one sector and the other can be 127 sectors, both could be 64 sectors, or any other combination as long as the total is a multiple of 128. In the following example, there is one subdisk per column.

To create an LSM volume that uses a striped plex on different buses:

1. Create the subdisks on disks on different buses:

   ```
   # volmake [-g disk_group] sd subdisk disk len=length
   ```

The following example creates subdisks on disks dsk2 and dsk3 (on bus 1), dsk4 and dsk5 (on bus 2), and dsk6 and dsk7 (on bus 3):

```
# volmake sd dsk2-01 dsk2 len=16m
# volmake sd dsk3-01 dsk3 len=16m
# volmake sd dsk4-01 dsk4 len=16m
# volmake sd dsk5-01 dsk5 len=16m
# volmake sd dsk6-01 dsk6 len=16m
# volmake sd dsk7-01 dsk7 len=16m
```

2. Create a striped plex, specifying the order of subdisks on which to create the columns:

```
# volmake [-g disk_group] plex plex layout=stripe \
stwidth=data_unit_size sd=subdisk,...
```

The following example uses the subdisks created in step 1 and lists them in alternating bus order to create a six-column striped plex called plex_01. The command line lists the subdisks in a pattern that alternates the bus order from bus 1 to 2, then bus 1 to 3, then bus 2 to 3:

```
# volmake plex plex_01 layout=stripe stwidth=64 \
sd=dsk2-01,dsk4-01,dsk3-01,dsk6-01,dsk5-01,dsk7-01
```

3. If you want to create a mirror plex and a DRL for the volume, complete this step. If the volume will have only one data plex, go to step 4.

   a. Repeat step 1 to create subdisks on a different group of disks on different buses for the second data plex.

      The following example creates subdisks for the columns in the second data plex on disks dsk8 and dsk9 (on bus 4), dsk10 and dsk11 (on bus 5), and dsk12 and dsk13 (on bus 6):

```
# volmake sd dsk8-01 dsk8 len=16m
# volmake sd dsk9-01 dsk9 len=16m
# volmake sd dsk10-01 dsk10 len=16m
# volmake sd dsk11-01 dsk11 len=16m
# volmake sd dsk12-01 dsk12 len=16m
# volmake sd dsk13-01 dsk13 len=16m
```

   b. Repeat step 2 to create the second data plex, specifying the order of subdisks on which to create the columns.

      The following example uses the subdisks created in step 3a and lists them in alternating bus order to create a six-column striped plex called plex_02. The command line lists the subdisks in a pattern that alternates the bus order from bus 4 to 5, then bus 4 to 6, then bus 5 to 6:

```
# volmake plex plex_02 layout=stripe stwidth=64 \
sd=dsk8-01,dsk10-01,dsk9-01,dsk12-01,dsk11-01,dsk13-01
```

4. Create the LSM volume, specifying the name of the data plex you created in step 2, and the additional data plex (if any) you created in step 3:

   ```
   # volmake [-g disk_group] -U usage_type vol volume \
   plex=plex ...
   ```

   The following example creates an LSM volume called vol9 with a usage type of `fsgen`, using a plex called plex_stripe:

   ```
   # volmake -U fsgen vol vol9 plex=plex_stripe
   ```

   The following example creates an LSM volume called vol_mirr with a usage type of `fsgen`, using data plexes called plex_01 and plex_02:

   ```
   # volmake -U fsgen vol vol_mirr plex=plex_01,plex_02
   ```

5. If the volume has mirror plexes, add a DRL plex to the volume on a disk that is not used by one of the data plexes:

   ```
   # volassist addlog volume disk
   ```

6. Start the LSM volume:

   ```
   # volume start volume
   ```

   The volume is ready for use.

4.2.4 Creating an LSM Volume Using a RAID 5 Plex

By default, the `volassist` command creates columns for a RAID 5 plex on disks in alphanumeric order, regardless of their order on the command line.

To improve performance, you might want to create the columns on disks on different buses. See Section 4.2.5 for more information about specifying the disk order for columns in a RAID 5 plex.

The `volassist` command automatically creates a RAID 5 log plex for the volume.

To create an LSM volume that uses a RAID 5 plex, enter:

```
# volassist [-g disk_group] make volume length layout=raid5 \
[nstripe=number_of_columns] [stripeunit=data_unit_size] [disks]
```

The following example creates a 6 GB, six-column volume called vol6 in a disk group called dg1, using any available disks:

```
# volassist -g dg1 make vol6 6g layout=raid5 nstripe=6
```

4.2.5 Creating an LSM Volume Using a RAID 5 Plex (on Different Buses)

By default, LSM creates columns for a RAID 5 plex on the first available disks it finds in the disk group. This might result in a volume with columns that use disks on the same bus.

You can improve performance by creating a RAID 5 plex with columns that use disks on different buses. To do so, you must create the subdisks for each column.

Note

Each column of subdisks should be the same size and be a multiple of the data unit size so there is no wasted space. For example, a data unit size (stripe width) of 16 KB for a RAID 5 plex corresponds to 32 blocks (sectors), so the total of the subdisks in each column should be a multiple of 32.

If each column comprises one subdisk (the typical configuration), then the subdisk size should be a multiple of 32. If a column comprises two subdisks, one subdisk can be one sector and the other can be 31 sectors, both could be 16 sectors, or any other combination as long as the total is a multiple of 32. In the following example, there is one subdisk per column.

To create an LSM volume that uses a RAID 5 plex on different buses:

1. Create the subdisks on disks on different buses:

   ```
   # volmake [-g disk_group] sd subdisk disk,offset,length
   ```

   The following example creates 1 MB subdisks for the data plex on disks called dsk6, dsk7, dsk8, and dsk9. In this example, disks dsk6 and dsk7 are on bus 1, and dsk8 and dsk9 are on bus 2:

   ```
   # volmake sd dsk6-01 dsk6 len=1m
   # volmake sd dsk7-01 dsk7 len=1m
   # volmake sd dsk8-01 dsk8 len=1m
   # volmake sd dsk9-01 dsk9 len=1m
   ```

2. Create the RAID 5 data plex, specifying the order of subdisks on which to create the columns:

   ```
   # volmake [-g disk_group] plex plex layout=raid5 \
   stwidth=data_unit_size sd=subdisk,...
   ```

   The following example uses the subdisks created in step 1 to create a four-column RAID 5 data plex called plex–01:

```
# volmake plex plex-01 layout=raid5 stwidth=16 \
sd=dsk6-01,dsk8-01,dsk7-01,dsk9-01
```

Note that in this plex, the stripe alternates between subdisks on buses 1 and 2.

3. Create the LSM volume, specifying the data plex:

   ```
   # volmake [-g disk_group] -U raid5 vol volume plex=plex
   ```

   The following example creates an LSM volume called vol5 using the plex called plex–01:

   ```
   # volmake -U raid5 vol vol5 plex=plex-01
   ```

4. Add a RAID 5 log plex to the volume, on a disk that is not used by the data plex:

   ```
   # volassist addlog volume disk
   ```

5. Start the LSM volume:

   ```
   # volume start volume
   ```

   The volume is ready for use.

4.2.6 Creating an LSM Volume for Secondary Swap Space

If disk errors occur in the swap space, a system crash is likely to occur. You can create an LSM volume using mirrored concatenated plexes to protect against disk I/O errors in the secondary swap space. Do not create a DRL plex for swap volumes, because mirror resynchronization is not necessary, and a DRL plex on swap volumes will interfere with crash dumps.

To create an LSM volume for the secondary swap space:

1. Create an LSM volume without a log:

   ```
   # volassist [-g disk_group] -U gen make volume length \
   nmirror=n layout=nolog
   ```

   The following example creates an LSM volume called vol_swap2 that uses two mirrors with no log:

   ```
   # volassist -U gen make vol_swap2 128m nmirror=2 layout=nolog
   ```

2. Set the LSM volume with the `start_ops=norecov` option so LSM does not resynchronize the mirrors:

   ```
   # volume set start_opts=norecov volume
   ```

3. Add the LSM volume as secondary swap space using the `swapon` command:

   ```
   # swapon /dev/vol/volume
   ```

4. Add the LSM device special file to the `swapdevice` kernel attribute value within the `vm:` section of the `/etc/sysconfigtab` file. The following example shows the entry to change:

   ```
   vm:
   swapdevice=/dev/disk/dsk1b, /dev/vol/volume
   ```

See the *System Administration* manual and the `swapon`(8) and `sysconfig`(8) reference pages for more information on adding additional swap space.

4.3 Configuring File Systems to Use LSM Volumes

After you create an LSM volume, you use it the same way you would use a disk partition. Because LSM uses the same interfaces as disk device drivers, you can specify an LSM volume in any operation where you can specify a disk or disk partition.

The following sections describe how to configure AdvFS and UFS to use an LSM volume.

4.3.1 Creating an AdvFS Domain on an LSM Volume

AdvFS treats LSM volumes as it does any other storage device. See *AdvFS Administration* for information on creating an AdvFS domain.

Note

> You can add more LSM volumes to an existing AdvFS domain if the domain needs more storage by creating a new LSM volume and using the AdvFS `addvol` command to add the volume to the domain. See *AdvFS Administration* for more information.

4.3.2 Creating a UFS File System on an LSM Volume

To create a UFS on an LSM volume:

1. Create a UFS using the LSM disk group and volume name:

   ```
   # newfs [options] /dev/rvol/disk_group/volume
   ```

   The following example creates a UFS on an LSM volume called vol_ufs in the dg1 disk group:

   ```
   # newfs /dev/rvol/dg1/vol_ufs
   ```

   It is not necessary to specify the name of the disk group for LSM volumes in the rootdg disk group.

   See the `newfs`(8) reference page for information on `newfs` options.

2. Use the LSM block special device name to mount the file system:

```
# mount /dev/vol/disk_group/volume /mount_point
```

The following example mounts the LSM volume called vol_ufs as mnt2:

```
# mount /dev/vol/dg1/vol_ufs /mnt2
```

4.4 Creating an LSM Volume for Existing Data

When you create an LSM volume for existing data on a disk or partition, LSM:

- Converts the disk or partition to an LSM nopriv disk
- Encapsulates the data in the disk or disk partition
- Configures the disk or disk partition into an LSM volume in the rootdg disk group

You can encapsulate data in:

- Disks or disk partitions, including UFS file systems (Section 4.4.1)
- AdvFS domains (Section 4.4.2)
- The boot disk (Section 3.4.3)

4.4.1 Creating a Volume from Disks or Disk Partitions

The encapsulation procedure configures disks and disk partitions, which can contain any kind of data including a UFS file system, into LSM nopriv disks using information in the disk label and in the `/etc/fstab` file. After the encapsulation, entries in the `/etc/fstab` file or in the `/etc/sysconfigtab` file are changed to use the LSM volume name instead of the block device name of the disk or disk partition.

If you encapsulate an entire disk (by not specifying a partition letter), such as dsk3, all the in-use partitions are encapsulated as one LSM nopriv disk.

To encapsulate a disk or disk partition:

1. Back up the data on the disk or disk partition to be encapsulated.
2. Unmount the disk or partition or take the data off line. If you cannot unmount the disk or partition or take the data off line, you must restart the system to complete the encapsulation procedure.
3. Create the LSM encapsulation script:

   ```
   # volencap [-g disk_group] {disk|partition}
   ```

   The following example creates an encapsulation script for a disk called dsk3:

   ```
   # volencap dsk3
   ```

Note

Although you can encapsulate several disks or disk partitions at the same time, it is recommended that you encapsulate each disk or disk partition separately.

4. Complete the encapsulation process:

```
# volreconfig
```

If the encapsulated disk or disk partition is in use, the `volreconfig` command prompts you to restart the system.

5. Optional but recommended, move the volume off the nopriv disks to simple or sliced disks in the same disk group. See Section 5.1.6 for more information.

4.4.2 Creating a Volume from an AdvFS Domain

You can place an existing AdvFS domain into an LSM volume by either encapsulating the domain or migrating the domain to an LSM volume:

- Encapsulating the domain creates an LSM volume on the same disk or disks that the domain already uses.
  - If an AdvFS domain consists of one disk or partition, you can encapsulate the disk or partition using the procedure described in Section 4.4.1.
  - If the AdvFS domain consists of multiple disks or partitions (requires the AdvFS Utilities License), you can encapsulate the AdvFS domain instead of the individual disk or partitions.

  LSM creates an LSM volume for each AdvFS disk or partition in the domain.

  Encapsulating a domain might require restarting the system if you cannot unmount the filesets before performing the encapsulation.
- Migrating the domain creates an LSM volume on disks that you specify, moves the domain data to the new volume, and removes the domain from its original disks. The disks are no longer in use by the domain after the migration completes.

  This operation does not require you to unmount filesets or restart the system, but temporarily uses additional disk space until the migration is complete.

When you encapsulate an AdvFS domain, LSM changes the links in the `/etc/fdmns` directory to point to the LSM volumes.

No mount point changes are necessary during encapsulation or migration, because the mounted filesets are abstractions to the domain. The domain can be activated normally after the encapsulation or migration process completes. Once the domain is activated, the filesets remain unchanged and the encapsulation or migration is transparent to AdvFS domain users.

4.4.2.1 Encapsulating an AdvFS Domain

To encapsulate an AdvFS domain:

1. Back up the data in the AdvFS domain with the `vdump` utility.
2. Unmount all filesets.

   If the domain is in use (you cannot unmount the filesets), you can create the encapsulation script and run `volreconfig` when convenient to complete the encapsulation procedure.
3. Create the LSM encapsulation script:

   ```
   # volencap domain
   ```

   The following example creates an encapsulation script for an AdvFS domain called dom1:

   ```
   # volencap dom1
   ```
4. Complete the encapsulation procedure:

   ```
   # volreconfig
   ```

   If the AdvFS domain is mounted, the `volreconfig` command prompts you to restart the system.

   The `/etc/fdmns` directory is updated on successful creation of LSM volumes.

4.4.2.2 Migrating an AdvFS Domain

You can place any AdvFS domain (except for the `root_domain`) into an LSM volume.

Note on TruClusters

See the *Cluster Administration* manual for information on migrating an AdvFS domain in a cluster.

This operation uses a different disk than the disk on which the domain originally resides and therefore does not require a restart. You can specify:

- The name of the volume (default is the name of the domain with the suffix `vol`)

- The number of stripes and mirrors that you want the volumes to use

  Striping improves read performance, and mirroring ensures data availability in the event of a disk failure.

You must specify LSM disks by their disk media names to create the volume for the domain. The disks that you specify must belong to the same disk group, because LSM volumes can use disks from only one disk group. There must be sufficient LSM disks and the disks must be large enough to contain the domain. See the `volmigrate(8)` reference page for more information on disk requirements and the options for striping and mirroring.

To migrate a domain into an LSM volume, enter:

```
# volmigrate [-g diskgroup] [-m num_mirrors] [-s num_stripes]
domain disk_media_name ...
```

The `volmigrate` command creates a volume with the specified characteristics, moves the data from the domain into the volume, removes the original disk or disks from the domain, and leaves those disks unused. The volume is started and ready for use, and no restart is required.

5 Managing LSM Objects

This chapter describes how to manage LSM objects using LSM commands. You can also accomplish the tasks described in this chapter using:

- The Storage Administrator (Appendix A)
- The `voldiskadm` menu interface (Appendix C)
- The Visual Administrator (Appendix D)

For more information on an LSM command, see the reference page corresponding to its name. For example, for more information on the `volassist` command, enter:

```
# man volassist
```

5.1 Managing LSM Disks

The following sections describe how to use LSM commands to manage LSM disks.

5.1.1 Creating an LSM Disk

You create an LSM disk when you initialize a disk or partition for LSM use. When you initialize a disk or partition for LSM use, LSM:

- Destroys data on the disk
- Formats the disk as an LSM disk
- Assigns it a disk media name
- Writes a new disk label

You can configure an LSM disk in a disk group or as a spare disk. If you configure the LSM disk in a disk group, LSM uses it to store data. If you configure an LSM disk as a spare, LSM uses it as a replacement for a failed LSM disk that contains a mirror or RAID 5 plex.

If the disk is new to the system, enter the `voldctl enable` command after entering the `hwmgr -scan scsi` command to make LSM recognize the disk.

To initialize a disk or partition as an LSM disk, you can use:

- The `voldiskadd` script as described in Section 4.1.2.

- The voldisksetup command.

  # **voldisksetup -i {*disk*|*partition*}**

  Note

  By default, LSM initializes each disk with one copy of the configuration database. If a disk group will have fewer than four disks, you should initialize each disk to have two copies of the disk group's configuration database to ensure that the disk group has multiple copies in case one or more disks fail. You must use the voldisksetup command to enable multiple copies of the configuration database on a disk.

  Specifying a disk access name initializes the entire disk as an LSM sliced disk. Specifying a partition name initializes the partition as an LSM simple disk.

  To initialize one or more disks, optionally setting the number of configuration copies to 2:

  # **voldisksetup -i *disk* ...** [**nconfig=2**]

  After you initialize a disk or disk partition as an LSM disk, you can add it to a disk group. See Section 5.2.2 for information on creating a disk group or Section 5.2.3 for information on adding an LSM disk to an existing disk group.

5.1.2 Displaying LSM Disk Information

To display detailed information for an LSM disk, enter:

voldisk list *disk*

The following example contains information for an LSM disk called dsk5:

```
Device:    dsk5
devicetag: dsk5
type:      sliced
hostid:    servername
disk:      name=dsk5 id=942260116.1188.servername
group:     name=dg1 id=951155418.1233.servername
flags:     online ready autoimport imported
pubpaths:  block=/dev/disk/dsk5g char=/dev/rdisk/dsk5g
privpaths: block=/dev/disk/dsk5h char=/dev/rdisk/dsk5h
version:   n.n
iosize:    min=512 (bytes) max=2048 (blocks)
public:    slice=6 offset=16 len=2046748
private:   slice=7 offset=0 len=4096
update:    time=952956192 seqno=0.11
headers:   0 248
configs:   count=1 len=2993
logs:      count=1 len=453
Defined regions:
 config   priv     17-   247[   231]: copy=01 offset=000000 enabled
```

```
config   priv    249-  3010[  2762]: copy=01 offset=000231 enabled
log      priv   3011-  3463[   453]: copy=01 offset=000000 enabled
```

5.1.3 Renaming an LSM Disk

When you initialize an LSM disk, you can assign it a disk media name or use the default disk media name, which is the same as the disk access name assigned by the operating system software.

Caution

Each disk in a disk group must have a unique name. To avoid confusion, you might want to ensure that no two disk groups contain disks with the same name. For example, both the rootdg disk group and another disk group could contain disks with a disk media name of disk03. Because most LSM commands operate on the rootdg disk group unless you specify otherwise, you might perform operations on the wrong disk if multiple disk groups contain identically named disks.

The `voldisk list` command displays a list of all the LSM disks in all disk groups on the system.

To rename an LSM disk, enter:

```
# voledit rename old_disk_media_name new_disk_media_name
```

For example, to rename an LSM disk called disk03 to disk01, enter:

```
# voledit rename disk03 disk01
```

5.1.4 Placing an LSM Disk Off Line

You can place an LSM disk off line to:

- Prevent LSM from accessing it
- Enable you to move the disk to a different physical location and have the disk retain its LSM identity

Placing a disk off line closes its device file. You cannot place an LSM disk off line if it is in use.

To place an LSM disk off line:

1. Remove the LSM disk from its disk group:

   ```
   # voldg -g disk_group rmdisk disk
   ```

2. Place the LSM disk off line:

   ```
   # voldisk offline disk
   ```

5.1.5 Placing an LSM Disk On Line

To restore access to an LSM disk that you placed off line, you must place it on line. The LSM disk is placed in the free disk pool and is accessible to LSM again. After placing an LSM disk on line, you must add it to a disk group before an LSM volume can use it. If the disk belonged to a disk group previously, you can add it to the same disk group.

To place an LSM disk on line, enter:

```
# voldisk online disk
```

See Section 5.2.3 for information on adding an LSM disk to a disk group.

5.1.6 Moving Data from an LSM Disk

You can move (evacuate) LSM volume data to other LSM disks in the same disk group if there is sufficient free space. If you do not specify a target LSM disk, LSM uses any available LSM disk in the disk group that has sufficient free space. Moving data off an LSM disk is useful in the event of disk failure or to move a volume to simple or sliced disks after encapsulating a disk or partition.

Note

Do not move the contents of an LSM disk to another LSM disk that contains data from the same volume. If the volume is redundant (uses mirror plexes or a RAID 5 plex), the resulting layout might not preserve the volume's redundancy.

To move data off an LSM disk, enter:

```
# volevac [-g disk_group] source_disk target_disk
```

For example, to move data in the rootdg disk group from LSM disk dsk8 and to dsk9, enter:

```
# volevac dsk8 dsk9
```

5.1.7 Removing an LSM Disk from LSM Control

You can remove a disk from LSM control if you removed the disk from its disk group or deported its disk group.

See Section 5.2.8 for information on removing an LSM disk from a disk group. See Section 5.2.5 for information on deporting a disk group.

To remove an LSM disk, enter:

```
# voldisk rm disk
```

For example, to remove an LSM disk called dsk8, enter:

```
# voldisk rm dsk8
```

If you want to use the disk after it is removed from LSM control, you must initialize it using the `disklabel` command. See the `disklabel`(8) reference page for more information.

5.2 Managing Disk Groups

The following sections describe how to use LSM commands to manage disk groups.

5.2.1 Displaying Disk Group Information

There are three common ways to display disk group information. You can display:

- A list of all disks on the system. See Section 5.2.1.1.
- A list of disks in all disk groups and the free space on each. See Section 5.2.1.2.
- The maximum size volume you can create in a disk group. See Section 5.2.1.3.

5.2.1.1 Displaying LSM Disks in All Disk Groups

To display a list of all LSM disks and the disk group to which each belongs, enter:

```
# voldisk list
```

Information similar to the following is displayed:

```
DEVICE       TYPE      DISK         GROUP        STATUS
dsk0         sliced    -            -            unknown
dsk1         sliced    -            -            unknown
dsk2         sliced    dsk2         rootdg       online
dsk3         sliced    dsk3         rootdg       online
dsk4         sliced    dsk4         rootdg       online
dsk5         sliced    dsk5         rootdg       online
dsk6         sliced    dsk6         dg1          online
dsk7         sliced    dsk7         dg1          online
dsk8         sliced    dsk8         dg1          online
dsk9         sliced    dsk9         dg2          online
dsk10        sliced    dsk10        dg2          online
dsk11        sliced    dsk11        dg2          online
dsk12        sliced    -            -            unknown
dsk13        sliced    -            -            unknown
```

The following list describes the preceding information categories:

DEVICE	The disk access name assigned by the operating system software.
TYPE	The LSM disk type: sliced, simple, or nopriv.
DISK	The LSM disk media name. An LSM disk media name is displayed only if the disk is in a disk group.
GROUP	The disk group to which the disk belongs. A disk group name is displayed only if the disk is in a disk group.
STATUS	The status of the LSM disk: • `online` — The disk is detected by LSM and is running. • `offline` — The disk has not been detected or was put off line manually. • `unknown` — The disk was detected but is not initialized for use by LSM. • `error` — The disk is detected but has experienced I/O errors. • `failed was` — An LSM disk media name exists but the disk is not associated with a `DEVICE`. Displays the last device associated with this name.

5.2.1.2 Displaying Free Space in Disk Groups

To display the free space in one or all disk groups, enter:

voldg [**-g** ***disk_group***] **free**

Information similar to the following is displayed:

```
GROUP       DISK        DEVICE      TAG         OFFSET    LENGTH    FLAGS
rootdg      dsk2        dsk2        dsk2        2097217   2009151   -
rootdg      dsk3        dsk3        dsk3        2097152   2009216   -
rootdg      dsk4        dsk4        dsk4        0         4106368   -
rootdg      dsk5        dsk5        dsk5        0         4106368   -
dg1         dsk6        dsk6        dsk6        0         2046748   -
dg1         dsk8        dsk8        dsk8        0         2046748   -
```

The value in the `LENGTH` column indicates the amount of free disk space in 512–byte blocks. (2048 blocks equal 1 MB.)

5.2.1.3 Displaying the Maximum Size for an LSM Volume in a Disk Group

To display the maximum size for an LSM volume that you can create in a disk group, enter:

volassist [**-g** ***disk_group***] **maxsize**

The following example displays the maximum size for an LSM volume that you can create in a disk group called dg1:

```
# volassist -g dg1 maxsize
Maximum volume size: 6139904 (2998Mb)
```

5.2.2 Creating a Disk Group

The default rootdg disk group is created when you install LSM and always exists on a system running LSM. You can create additional disk groups to organize your disks into logical sets. Each disk group that you create must have a unique name and contain at least one simple or sliced LSM disk. An LSM disk can belong to only one disk group. An LSM volume can use disks from only one disk group.

If you want to initialize LSM disks and create a new disk group at the same time, you can use the `voldiskadd` script. (See Section 4.1.2 for more information.)

Note

By default, LSM initializes each disk with one copy of the configuration database. If a disk group will have fewer than four disks, you should initialize each disk to have two copies of the disk group's configuration database to ensure that the disk group has multiple copies in case one or more disks fail. You must use the `voldisksetup` command to enable more than one copy of the configuration database (Section 5.1.1).

To create a new disk group using LSM disks, enter:

```
# voldg init new_disk_group disk ...
```

For example, to create a disk group called newdg using LSM disks called dsk100, dsk101, and dsk102, enter:

```
# voldg init newdg dsk100 dsk101 dsk102
```

5.2.3 Adding a Disk to a Disk Group

To add a disk to an existing disk group, enter:

```
# voldg [-g disk_group] adddisk disk
```

For example, to add the disk called dsk10 to a disk group called dg1, enter:

```
# voldg -g dg1 adddisk dsk10
```

5.2.4 Renaming a Disk Group

Renaming a disk group involves deporting and then importing the disk group. You cannot rename a disk group while it is in use. All activity on all volumes in the disk group must stop, and the volumes in the disk group are inaccessible while the disk group is deported.

Because renaming a disk group involves an interruption of service to the volumes, you might want to perform this task during a planned shutdown or maintenance period. Choose the new disk group name carefully, and ensure that the new name is easy to remember and use. Renaming a disk group updates the `/etc/fstab` file.

Note

You cannot rename the rootdg disk group.

To rename a disk group:

1. Deport the disk group, assigning it a new name. See Section 5.2.5.
2. Import the disk group using its new name. See Section 5.2.6.

5.2.5 Deporting a Disk Group

Deporting a disk group makes its volumes inaccessible. You can deport a disk group to:

- Rename the disk group.
- Reuse the disks for other purposes.
- Move the disk group to another system (Section 5.2.7).

You cannot deport the rootdg disk group.

Caution

The `voldisk list` command displays the disks in a deported disk group as available (with a status of `online`). However, removing or reusing the disks in a deported disk group can result in data loss.

To deport a disk group:

1. If volumes in the disk group are in use, stop the volumes:

   ```
   # volume [-g disk_group] stopall
   ```

2. Deport the disk group:

   - To deport the disk group with no changes, enter:

     ```
     # voldg deport disk_group
     ```

   - To deport the disk group and assign it a new name, enter:

```
# voldg [-n newname] deport disk_group
```

See the voldg(8) reference page for more information on assigning a new name to a disk group.

You must import a disk group before you can use it. See Section 5.2.6 for information on importing a disk group.

If you no longer need the disk group, you can:

- Add the disks to different disk groups (Section 5.2.3).
- Use the disks to create new disk groups (Section 5.2.2).
- Remove the disks from LSM control (Section 5.1.7).

5.2.6 Importing a Disk Group

Importing a disk group makes the disk group and its volumes accessible. You cannot import a disk group if you used any of its associated disks while it was deported.

To import a disk group and restart its volumes:

1. Import the disk group:

   ```
   # voldg import disk_group
   ```

2. Start all volumes within the disk group:

   ```
   # volume [-g disk_group] startall
   ```

5.2.7 Moving a Disk Group to Another System

You can move a set of disks from one system to another and retain the LSM objects and data on those disks. You can move any disk group to another system; however, to move the rootdg disk group, the following must be true:

- You have stopped LSM running on the original system. (You stopped the LSM daemons or shut down the system before disconnecting the rootdg disk group to move it.)
- The new system must be running LSM, which means it has a rootdg disk group; therefore, you must import the former rootdg disk group with a different name on the new system.

Moving a disk group between systems results in the new host system assigning new disk access names to the disks. For LSM nopriv disks (created when you encapsulate disks or partitions), the association between the original disk access name and its disk media name might be lost, or might be reassociated incorrectly. To prevent this, you must manually reassociate the disk media names with the new disk access names. For LSM simple and sliced disks, LSM manages this reassociation.

If possible, before moving the disk group, migrate the data from nopriv disks to simple or sliced disks, which have a private region and will be reassociated automatically. See Section 5.1.6 for more information on moving data to a different disk.

If you cannot move the data to simple or sliced disks, follow these steps to help ensure you can correctly reassociate the nopriv disks on the new host system:

1. On the original system, identify all the nopriv disks in the disk group by their current disk access name, disk media name, and a unique identifier (such as the disk's SCSI world-wide identifier) that will not change or can be tracked when the disk is connected to the new system.

   If there is only one nopriv disk in the disk group, there is only one device to reassociate. As long as you are not connecting other devices to the new host at the same time, you might not need this information. For two or more noprivs, having precise identification beforehand is crucial.
2. Keep track of the before-and-after bus locations of each nopriv disk as you move it between systems. Then when you scan for the disks on the new host, you will know which new disk access name associated with the new bus location belongs to which disk media name. You can move each disk individually and have the new host scan for it each time to be sure.

You can change the disk group's name or host ID when you move it to the new host:

- If the disk group's name is similar to a disk group on the system receiving it, you can change its name to reduce the chance for confusion.

  Note

  If the disk group has the same name as a disk group on the system receiving it, you must change its name.

- If you want the system receiving the disk group to import it automatically the first time the system starts up, you can change the disk group's host ID to that of the receiving system as you deport it from the original system.

  If you will import the disk group to a system that is already running, you do not need to change the disk group's host ID; it is changed as the disk group is imported.

To move a disk group to another system:

1. Stop all activity on the volumes in the disk group and unmount any file systems.

2. Deport the disk group from the originating system:

   - To deport the disk group with no changes, enter:

     ```
     # voldg deport disk_group
     ```

   - To deport the disk group and assign it a new host ID or a new name, enter:

     ```
     # voldg [-n newname] [-h newhostID] deport disk_group
     ```

3. Physically move the disks to the new host system.

4. Enter the following command on the new host system to scan for the disks:

   ```
   # hwmgr -scan scsi
   ```

   The `hwmgr` command returns the prompt before it completes the scan. You need to know that the system has discovered the disks before continuing. See the `hwmgr(8)` reference page for more information on how to trap the end of a scan.

5. Make the `vold` daemon scan for the newly added disks:

   ```
   # voldctl enable
   ```

6. Import the disk group to the new host. If the disk group contains nopriv disks whose disk media names no longer correspond to their original disk access names, you might need to use the force (`-f`) option.

   ```
   # voldg [-f] import disk_group
   ```

7. If applicable, associate the disk media names for the nopriv disks to their new disk access names:

   ```
   # voldg -g disk_group -k adddisk
   disk_media_name=disk_access_name ...
   ```

8. Recover and start all startable volumes in the imported disk group. The following command performs any necessary recovery operations as a background task after starting the volumes.

   ```
   # volrecover -g disk_group -sb
   ```

9. Optionally, check for any detached plexes.

   ```
   # volinfo -p
   ```

   If the output lists any volumes as `Unstartable`, see Section 6.4.3 for information on how to proceed.

10. If necessary, start the remaining `Startable` volumes:

    ```
    # volume -g disk_group start volume1 volume2 ...
    ```

5.2.8 Removing an LSM Disk from a Disk Group

You can remove an LSM disk from a disk group; however, you cannot remove:

- The last disk in a disk group unless the disk group is deported. See Section 5.2.5 for information on deporting a disk group.
- Any disk that is in use (for example, disks that contain active LSM volume data). If you attempt to remove a disk that is in use, LSM displays an error message and does not remove the disk.

  See Section 5.1.6 for information on moving data from a disk. See Section 5.4.6 for information on removing LSM volumes.

To remove an LSM disk from a disk group:

1. Verify that the LSM disk is not in use by listing all subdisks:

   `# ` **`volprint -st`**

   Information similar to the following is displayed:

   ```
   Disk group: rootdg

   SD NAME         PLEX         DISK     DISKOFFS LENGTH   [COL/]OFF DEVICE   MODE

   sd dsk1-01      klavol-01    dsk1     0        1408     0/0       dsk1     ENA
   sd dsk2-02      klavol-03    dsk2     0        65       LOG       dsk2     ENA
   sd dsk2-01      klavol-01    dsk2     65       1408     1/0       dsk2     ENA
   sd dsk3-01      klavol-01    dsk3     0        1408     2/0       dsk3     ENA
   sd dsk4-01      klavol-02    dsk4     0        1408     0/0       dsk4     ENA
   sd dsk5-01      klavol-02    dsk5     0        1408     1/0       dsk5     ENA
   sd dsk6-01      klavol-02    dsk6     0        1408     2/0       dsk6     ENA
   ```

   The disks in the `DISK` column are currently in use by LSM volumes, and therefore you cannot remove those disks from a disk group.

2. Remove the LSM disk from the disk group:

   `# ` **`voldg [-g`** ***`disk_group`*****`] rmdisk`** ***`disk`***

   For example, to remove an LSM disk called dsk8 from the rootdg disk group, enter:

   `# ` **`voldg rmdisk dsk8`**

The disk remains under LSM control. You can:

- Add the disk to a different disk group. See Section 5.2.3.
- Use the disk to create a new disk group. See Section 5.2.2.
- Remove the disk from LSM control. See Section 5.1.7.

5.3 Managing the LSM Configuration Database

This section describes how to manage the LSM configuration database, including:

- Backing up the configuration database
- Restoring the configuration database from backup
- Modifying the configuration database properties

5.3.1 Backing Up the LSM Configuration Database

One important responsibility in managing a system with LSM is to periodically make a backup copy of the LSM configuration database. This helps you:

- Restore volumes from backup after a major system failure
- Recreate LSM volumes after disk failures, if the failures resulted in the loss of all configuration database copies

The saved configuration database (also called a description set) is a record of the objects in the LSM configuration (the LSM disks, subdisks, plexes and volumes) and the disk group to which each object belongs.

Whenever you make a change to the LSM configuration, the backup copy becomes obsolete. As with any backup, the content is useful only as long as it accurately represents the current information. Any time the number, nature, or name of LSM objects change, consider making a backup of the LSM configuration database. The following list describes some of the changes that will invalidate a configuration database backup:

- Creating disk groups
- Adding or removing disks from disk groups or from LSM control
- Creating or removing volumes
- Changing the properties of volumes, such as the plex layout or number of logs

Note

Backing up the configuration database does not save the data in the volumes and does not save the configuration data for any volumes associated with the boot disk, if you encapsulated the boot disk.

Depending on the nature of a boot disk failure, you might need to restore the system partitions from backups or installation media to return to a state where the system partitions are not under LSM control. From there, you can redo the procedures to encapsulate the boot disk partitions into LSM volumes and add mirror plexes to those volumes.

See Section 6.5.6 for more information about recovering from a boot disk failure under LSM control.

See Section 5.4.2 for information on backing up volume data.

By default, LSM saves the entire configuration database to a time-stamped directory called `/usr/var/lsm/db/LSM.`*`date.hostname`*. You can specify a different location for the backup, but the directory must not exist.

In the directory, the backup procedure creates:

- A file called `header`, which contains host ID and checksum information, and a list of the other files in this directory.
- A copy of the `volboot` file.
- A file called `voldisk.list`, which contains a list of all LSM disks, their type (sliced, simple, nopriv), the size of their private and public regions, their disk group, and other information.
- A subdirectory called `rootdg.d`, which contains the `allvol.DF` file. The `allvol.DF` file contains detailed descriptions of every LSM subdisk, plex, and volume, describing all their properties and attributes.

To back up the LSM configuration database:

1. Enter the following command, optionally specifying a directory location other than the default to store the LSM configuration database:

   `#` **`volsave`** [**`-d`** ***`directory`***]

2. Save the backup to tape or other removable media.

You can save multiple versions of the configuration database; each new backup is saved in the `/usr/var/lsm/db` directory with its own date and time stamp, as shown in the following example:

```
dr-xr-x---   3 root     system      8192 May  5 09:36 LSM.20000505093612.hostname
dr-xr-x---   3 root     system      8192 May 10 10:53 LSM.20000510105256.hostname
```

5.3.2 Restoring the LSM Configuration Database from Backup

You use the `volrestore` command to restore an LSM configuration database that you saved with the `volsave` command. You can restore the configuration database of a specific disk group or volume or the entire configuration (all disk groups and volumes except those associated with the boot disk). If you have saved multiple versions of the configuration, you can choose a specific one to restore. If you do not choose one, LSM restores the most recent version.

Notes

Restoring the configuration database does not restore data in the LSM volumes. See Section 5.4.3 for information on restoring volume data.

The `volrestore` command does not restore volumes associated with the root (`/`), `/usr`, and `/var` file systems and the primary swap area. If volumes for these partitions are corrupted or destroyed, these partitions must be reencapsulated to use LSM volumes.

See the *Cluster Administration* manual for information about using the `volrestore` command in a cluster.

To restore a backed-up LSM configuration database:

1. Optionally, display a list of all available database backups:

   ```
   # ls /usr/var/lsm/db
   ```

   If you saved the configuration database to a different directory, specify that directory.

2. Restore the chosen configuration database:

   - To restore the entire configuration database, enter:

     ```
     # volrestore [-d directory]
     ```

   - To restore a specific disk group configuration database, enter:

     ```
     # volrestore [-d directory] -g disk_group
     ```

   - To restore a specific volume configuration database, enter:

     ```
     # volrestore [-d directory] -v volume
     ```

   - To restore a configuration database interactively, enabling you to select or skip specific objects, enter:

     ```
     # volrestore [-d directory] -i
     ```

3. Start the restored LSM volumes:

   ```
   # volume -g disk_group startall
   ```

   If the volumes will not start, you might need to manually edit the plex state. See Section 6.4.3.

4. If necessary, restore the volume contents (data) from backup. See Section 5.4.3 for more information.

5.3.3 Changing the Size and Number of Configuration Database Copies

LSM maintains copies of the configuration database on separate physical disks within each disk group. When the disk group runs out of space in the configuration database, LSM displays a message similar to the following:

```
volmake: No more space in disk group configuration
```

This might happen because:

- You imported an LSM configuration from a system running Tru64 UNIX Version 4.0, which used a smaller default private region size.
- One or more disks in the disk group contain two copies of the configuration database. Whenever a configuration change occurs, all active copies are updated. If one disk's copies cannot be updated because they have grown too large, then none of the copies for the whole disk group can be updated.

If the configuration database runs out of disk space and you determine that one or more disks have two copies of the configuration database, you can remove one copy from each disk that has two. However, make sure that there are sufficient copies of the configuration database available for redundancy. For example, if the disk group has a total of four copies and two are on the same disk, you should remove one copy from that disk and enable a copy on another disk that does not have any.

If all copies of the configuration database are the same size and no disk has more than one copy, this could indicate that the private regions of the disks are too small (for example, the disks were initialized on a system running an earlier version of LSM, with a smaller default private region). To resolve this problem, you must add new disks to LSM, which will have the larger default private region size, add the new disks to the disk group, and delete the copies of the configuration database on the other disks.

To reduce the number of configuration database copies:

1. Display information about the disk group's configuration database:

   ```
   # voldg list disk_group
   ```

   Example 5–1 shows output from the `voldg` command for a disk group with multiple copies of the configuration database on one disk. Example 5–2 shows output for a disk group from a previous version of the operating system, in which some disks have smaller private regions than the current default.

Example 5–1: Multiple Configuration Database Copies on a Disk

```
Group:   rootdg
dgid:   783105689.1025.lsm
import-id: 0.1
flags:
config:  seqno=0.1112 permlen=173 free=166 templen=6 loglen=26
config disk dsk13 copy 1 len=173 state=clean online
config disk dsk13 copy 2 len=173 state=clean online
config disk dsk11g copy 1 len=347 state=clean online
config disk dsk10g copy 1 len=347 state=clean online
log disk dsk11g copy 1 len=52
log disk dsk13 copy 1 len=26
log disk dsk13 copy 2 len=26
log disk dsk10g copy 1 len=52
```

In Example 5–1:

- The `len=` information in the lines beginning with `config disk` and `log disk` is the size in blocks of the amount of space available on the disk for copies of the configuration database or the log. The smallest length for a `config disk` or `log disk` limits the entire disk group by limiting the length of the configuration or log in memory.
- Disk dsk13 has two copies of the configuration database. This halves the total configuration space available in memory for the disk group and is therefore the limiting factor.

Example 5–2: Configuration Database Copy on a Disk with a Private Region Smaller Than the Current Default

```
Group:     rootdg
dgid:      921610896.1026.abc.xyz.com
import-id: 0.1  flags:
copies:    nconfig=default nlog=default
config:    seqno=0.1081 permlen=347 free=341 templen=3 loglen=52
config disk dsk7 copy 1 len=347 state=clean online
config disk dsk8 copy 1 len=2993 state=clean online
config disk dsk9 copy 1 len=2993 state=clean online
config disk dsk10 copy 1 len=2993 state=clean online
log disk dsk7 copy 1 len=52
log disk dsk8 copy 1 len=453
log disk dsk9 copy 1 len=453
log disk dsk10 copy 1 len=453
```

In Example 5–2:

- The `len=` information in the lines beginning with `config disk` and `log disk` is the size in blocks of the amount of space available on the disk for copies of the configuration database or the log. The smallest length for a `config disk` or `log disk` limits the entire disk group by limiting the length of the configuration or log in memory.

- Disk dsk7 has a smaller private region than the other disks, which means there is less space to store copies of the configuration database and log (in the line `config disk dsk7` there are 347 blocks available versus 2993 blocks on the other disks; in the line `log disk dsk7` there are 52 blocks available versus 453 blocks on the other disks). This restricts the disk group's ability to store additional records, because the smallest private region sets the limit for the group.

2. Modify the number of configuration copies in the disk group:
   - To reduce the number of copies on a disk, enter the following command where *n* is the number of copies you want the disk to retain:

     ```
     # voldisk moddb disk nconfig=n
     ```

     For example, to reduce the number of configuration copies on dsk13 from two to one, enter:

     ```
     # voldisk moddb dsk13 nconfig=1
     ```

   - To remove all copies from a disk, enter:

     ```
     # voldisk moddb disk nconfig=0
     ```

3. Optionally, display the new configuration by entering the following command:

   ```
   # voldg list rootdg
   ```

   Information similar to the following is displayed. In this example, the output shows the change to the configuration in Example 5–2:

   ```
   Group:     rootdg
   dgid:      921610896.1026.abc.xyz.com
   import-id: 0.1  flags:
   copies:    nconfig=default nlog=default
   config:    seqno=0.1091 permlen=2993 free=2987 templen=3 loglen=453
   config disk dsk7 copy 1 len=2993 state=clean online
   config disk dsk8 copy 1 len=2993 state=clean online
   config disk dsk9 copy 1 len=2993 state=clean online
   config disk dsk10 copy 1 len=2993 state=clean online
   log disk dsk7 copy 1 len=453
   log disk dsk8 copy 1 len=453
   log disk dsk9 copy 1 len=453
   log disk dsk10 copy 1 len=453
   ```

To add a copy to another disk to maintain the appropriate number of copies for the disk group:

1. Display a list of all disks in the disk group:

   ```
   # voldisk [-g disk_group] list
   ```

2. Compare the disks listed in the output of the `voldisk list` command to those listed in the output of the `voldg list` command to identify a disk in the disk group that does not have a copy of the configuration database.
3. Enable a copy on a disk that does not have one, using the disk access name:

   # **voldisk moddb *disk_access_name* nconfig=1**

5.4 Managing LSM Volumes

The following sections describe how to use LSM commands to manage LSM volumes. See Chapter 4 for information on creating LSM volumes.

5.4.1 Displaying LSM Volume Information

The `volprint` command displays information about LSM objects that make up an LSM volume. The following table lists the abbreviations used in `volprint` output:

Abbreviation	Specifies
`dg`	Disk group name
`dm`	Disk media name
`pl`	Plex name
`sd`	Subdisk name
`v`	LSM volume name

To display LSM object information for an LSM volume, enter:

volprint [**-g *disk_group***] **-ht *volume***

Information similar to the following is displayed:

```
Disk group: rootdg   1

V  NAME         USETYPE      KSTATE   STATE    LENGTH   READPOL   PREFPLEX
PL NAME         VOLUME       KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID MODE
SD NAME         PLEX         DISK     DISKOFFS LENGTH   [COL/]OFF DEVICE   MODE

v  klavol       fsgen        ENABLED  ACTIVE   4096     SELECT    -                2
pl klavol-01    klavol       ENABLED  ACTIVE   4224     STRIPE    3/128    RW      3
sd dsk1-01      klavol-01    dsk1     0        1408     0/0       dsk1     ENA     4
sd dsk2-01      klavol-01    dsk2     65       1408     1/0       dsk2     ENA
sd dsk3-01      klavol-01    dsk3     0        1408     2/0       dsk3     ENA
pl klavol-02    klavol       ENABLED  ACTIVE   4224     STRIPE    3/128    RW
sd dsk4-01      klavol-02    dsk4     0        1408     0/0       dsk4     ENA
sd dsk5-01      klavol-02    dsk5     0        1408     1/0       dsk5     ENA
sd dsk6-01      klavol-02    dsk6     0        1408     2/0       dsk6     ENA
pl klavol-03    klavol       ENABLED  ACTIVE   LOGONLY  CONCAT    -        RW
sd dsk2-02      klavol-03    dsk2     0        65       LOG       dsk2     ENA
```

This example shows output for a volume that uses a three-column, striped plex that has one mirror plex.

1 Disk group name.

2 Volume name (klavol), usage type (`fsgen`), state (ENABLED ACTIVE), and size (4096) information.

3 Plex information. This volume has two data plexes, klavol–01 and klavol–02, and a DRL plex, klavol–03.

4 Subdisk information for the plex klavol–01.

5.4.2 Backing Up an LSM Volume

One of the more common tasks of a system administrator is helping users recover lost or corrupted files. To perform that task effectively, you must set up procedures for backing up LSM volumes and the LSM configuration database at frequent and regular intervals. You will need the saved configuration database as well as the backed-up data if you need to restore a volume after a major failure. (For example, multiple disks in the same volume failed, and those disks contained the active configuration records for the disk group.)

See Section 5.3.1 for information on backing up the LSM configuration database.

Note

If the volume is part of an Advanced File System domain, use the `vdump` command to back up the volume. See *AdvFS Administration* for more information.

The way you back up an LSM volume depends on the number and type of plexes in the volume:

- If the volume has only one concatenated or striped plex, see Section 5.4.2.1.
- If the volume has mirror plexes, see Section 5.4.2.2.
- If the volume has a RAID 5 plex, see Section 5.4.2.3.

5.4.2.1 Backing Up a Volume with a Single Concatenated or Striped Plex

To back up an LSM volume that has a single plex:

1. If necessary, select a convenient time and inform users to save files and refrain from using the volume (the application or file system that uses the volume) while you back it up.

2. Determine the size of the LSM volume and which disks it uses:

```
# volprint -v [-g disk_group] volume
```

3. Ensure there is enough free space in the disk group to create a temporary copy of the LSM volume. The free space must be on disks that are not used in the volume you want to back up:

```
# voldg [-g disk_group] free
```

4. If the volume contains a UNIX File System, unmount it.

5. Create a temporary mirror plex for the LSM volume, running this operation in the background:

```
# volassist snapstart volume &
```

6. Create a new volume from the temporary plex. (The `snapshot` keyword automatically uses the temporary plex to create the new volume.)

```
# volassist snapshot volume temp_volume
```

The following example creates a temporary LSM volume called vol1_backup for an LSM volume called vol1:

```
# volassist snapshot vol1 vol1_backup
```

7. Remount and resume use of the original LSM volume.

8. Start the temporary LSM volume:

```
# volume start temp_volume
```

9. Back up the temporary LSM volume to your default backup device:

```
# dump 0 /dev/rvol/disk_group/temp_volume
```

The following example backs up an LSM volume called vol1_backup in the rootdg disk group:

```
# dump 0 /dev/rvol/rootdg/vol1_backup
```

10. Stop and remove the temporary LSM volume:

```
# volume stop temp_volume
# voledit -r rm temp_volume
```

See the `dump(8)` reference page for more information about the `dump` command.

5.4.2.2 Backing Up a Volume with Mirror Plexes

Volumes with mirror plexes can remain in use while you back up their data, but any writes to the volume during the backup might result in inconsistency between the volume's data and the data that was backed up.

Caution

If the LSM volume has only two data plexes, redundancy is not available during the backup.

To back up an LSM volume that has mirror plexes:

1. Dissociate one of the volume's plexes, which leaves the plex as an image of the LSM volume at the time of dissociation:

   ```
   # volplex dis plex
   ```

   The following example dissociates a plex called vol01–02:

   ```
   # volplex dis vol01-02
   ```

2. Create a temporary LSM volume using the dissociated plex:

   ```
   # volmake -U fsgen vol temp_volume plex=plex
   ```

   The following example creates an LSM volume called vol01–temp using a plex called vol01–02:

   ```
   #  volmake -U fsgen vol vol01-temp plex=vol01-02
   ```

3. Start the temporary volume:

   ```
   # volume start temp_volume
   ```

4. Back up the temporary LSM volume to your default backup device:

   ```
   # dump 0 /dev/rvol/disk_group/temp_volume
   ```

   The following example backs up an LSM volume called vol1_backup in the rootdg disk group:

   ```
   # dump 0 /dev/rvol/rootdg/vol1_backup
   ```

5. Stop and remove the temporary LSM volume:

   ```
   # volume stop temp_volume
   # voledit -r rm temp_volume
   ```

6. Reattach the dissociated plex to the original volume. If the volume is very large, you can run this operation in the background:

   ```
   # volplex att volume plex &
   ```

   LSM automatically resynchronizes the plexes when you reattach the dissociated plex. This operation might take a long time, depending on the size of the volume. Running this process in the background returns control of the system to you immediately instead of after the resynchronization is complete.

See the `dump(8)` reference page for more information about the `dump` command.

5.4.2.3 Backing Up a Volume with a RAID 5 Plex

You can back up a volume that uses a RAID 5 plex, but you must either stop all applications from using the volume while the backup is in process or allow the backup to occur while the volume is in use.

If the volume remains in use during the backup, the volume data might change before the backup completes, and therefore the backup data will not be an exact copy of the volume's contents.

To back up a volume with a RAID 5 plex, enter:

```
# dump 0 /dev/rvol/disk_group/volume
```

5.4.3 Restoring an LSM Volume from Backup

The way you restore an LSM volume depends on what the volume is used for and if the volume is configured and active.

Note

If the volume is part of an AdvFS domain, use the `vrestore` command to restore it. See *AdvFS Administration* for more information.

If the volume is used for an application such as a database, see that application's documentation for the recommended method for restoring backed-up data.

To restore a backed-up volume:

- If the volume contains a UNIX File System and the volume still exists (for example, you replaced a failed disk), enter the following command to restore the data from the backup media:

  ```
  # restore -Yf backup_volume
  ```

- If the volume does not exist:

  1. Recreate the volume:

     ```
     # volrestore [-g disk_group] -v volume
     ```

  2. Recreate the file system:

     ```
     # newfs /dev/rvol/disk_group/volume
     ```

  3. Mount the file system:

     ```
     # mount /dev/vol/disk_group/volume /mount_point
     ```

  4. Restore the volume data:

     ```
     # restore -Yrf backup_volume
     ```

5.4.4 Starting an LSM Volume

LSM automatically starts all startable volumes when the system boots. You can manually start an LSM volume that:

- You manually stopped
- Belongs to a disk group that you manually imported
- Stopped because of a disk failure or other problem that you have since resolved

To start an LSM volume, enter:

```
# volume start [-g disk_group] volume
```

To start all volumes in a disk group (for example, after importing the disk group), enter:

```
# volume [-g disk_group] startall
```

5.4.5 Stopping an LSM Volume

LSM automatically stops LSM volumes when the system shuts down. When you no longer need an LSM volume, you can stop it then remove it. You cannot stop an LSM volume if a file system is using it.

To stop an LSM volume:

1. If applicable, stop a file system from using the LSM volume.
   - For AdvFS, dissociate the volume from the domain:

     ```
     # rmvol LSM_volume domain
     ```

     Data on the volume is automatically migrated to other volumes in the domain, if available. See the *AdvFS Administration* manual for more information on the `rmvol` command.
   - For UFS, unmount the file system:

     ```
     # umount /dev/rvol/volume
     ```
2. Stop the LSM volume:

   ```
   # volume [-g disk_group] stop volume
   ```

   For example, to stop an LSM volume called vol1 in the dg1 disk group, enter:

   ```
   # volume -g dg1 stop vol1
   ```

   To stop all volumes, enter:

   ```
   # volume stopall
   ```

5.4.6 Removing an LSM Volume

Removing an LSM volume destroys the data in that volume. Remove an LSM volume only if you are sure that you do not need the data in the LSM volume or the data is backed up elsewhere. When an LSM volume is removed, the space it occupied is returned to the free space pool.

The following procedure also unencapsulates UNIX File Systems.

Note for AdvFS Domains

To remove a volume that was created by encapsulating an AdvFS domain, see Section 5.4.6.1.

To remove an LSM volume:

1. If applicable, stop a file system from using the LSM volume.
   - For AdvFS, dissociate the volume from the domain:

     ```
     # rmvol LSM_volume domain
     ```

     Data on the volume is automatically migrated to other volumes in the domain. See the *AdvFS Administration* manual for more information on the `rmvol` command.
   - For UFS, unmount the file system:

     ```
     # umount /dev/rvol/volume
     ```
2. Edit the necessary system files as follows:
   - If the volume was configured as secondary swap, remove references to the LSM volume from the `vm:swapdevice` entry in the `sysconfigtab` file.
   - If the swap space was configured using the `/etc/fstab` file, update this file to change the swap entries back to disk partitions instead of LSM volumes.

   These changes are effective the next time the system restarts.

   See the *System Administration* manual and the `swapon`(8) reference page for more information.
3. Stop the LSM volume:

   ```
   # volume [-g disk_group] stop volume
   ```
4. Remove the LSM volume:

   ```
   # voledit -r rm volume
   ```

   This step removes the plexes and subdisks and the volume itself.

5. If the volume contained an encapsulated UNIX file system, edit the /etc/fstab file to change the volume name to the disk name. For example, change /dev/vol/rootdg/vol-dsk4g to /dev/dsk4g.

5.4.6.1 Unencapsulating AdvFS Domains

When you encapsulate AdvFS domains into LSM volumes, LSM creates a script that you can run to unencapsulate the domain. The script contains some LSM commands and some general commands and performs all the steps necessary to remove the LSM volume, remove the disk from LSM control, and update the links in the /etc/fdmns directory.

The script is created in the following directory, where *dsknp* is the disk access name of the disk the domain resides on:

```
/etc/vol/reconfig.d/disk.d/dsknp.encapdone/recover.sh
```

If you have encapsulated more than one AdvFS domain, the /etc/vol/reconfig.d/disk.d directory will contain a subdirectory for each disk. Make sure you run the correct script to unencapsulate the domain.

Note

Unencapsulating an AdvFS domain requires that you unmount the filesets.

To unencapsulate an AdvFS domain:

1. Display and unmount all the filesets in the domain. For example, to unmount the filesets in the new_dom domain, enter:

```
# mount
root_domain#root on / type advfs (rw)
/proc on /proc type procfs (rw)
usr_domain#usr on /usr type advfs (rw)
var_domain#var on /var type advfs (rw)
mhs:/work on /work type nfs (v3, rw, udp, hard, intr)
new_dom#junk on /junk type advfs (rw)
new_dom#stuff on /stuff type advfs (rw)
# umount /junk /stuff
```

2. Identify the name of the LSM volume for the domain and the name of the disk the domain is using:

```
# showfdmn domain
```

Information similar to the following is displayed:

```
               Id              Date Created  LogPgs  Version  Domain Name
3a65b2a9.0004cb3f  Wed Jan 17 09:56:41 2001     512        4  new_dom

Vol   512-Blks        Free  % Used  Cmode  Rblks  Wblks  Vol Name
```

```
1L    8380080     8371248      0%     on    512    512  /dev/vol/rootdg/vol-dsk10c
```

Typically, the volume name is derived from the disk containing the encapsulated domain.

3. Stop the LSM volume:

   ```
   # volume stop volume
   ```

4. Run the unencapsulation script. For example, to run the script for the `new_dom` domain on disk dsk10c, enter:

   ```
   # sh /etc/vol/reconfig.d/disk.d/dsk10c.encapdone/recover.sh
   ```

5. If the script is not available, do the following:

   a. Change directory to the domain directory:

      ```
      # cd /etc/fdmns/domain
      ```

   b. Remove the link to the volume:

      ```
      # rm disk_group.volume
      ```

   c. Replace the link to the disk device file:

      ```
      # ln -s /dev/disk/dsknp
      ```

6. Remount the filesets to the domain:

   ```
   # mount new_dom#junk /junk
   # mount new_dom#stuff /stuff
   ```

The domain is available for use. I/O to the domain goes through the disk device path instead of the LSM volume. You can confirm this by running the `showfdmn` command again:

```
# showfdmn new_dom

               Id              Date Created  LogPgs  Version  Domain Name
3a65b2a9.0004cb3f  Wed Jan 17 09:56:41 2001     512        4  new_dom

  Vol   512-Blks        Free  % Used  Cmode  Rblks  Wblks  Vol Name
   1L    8380080     8371248      0%     on    256    256  /dev/disk/dsk10c
```

5.4.6.2 Unencapsulating System Partitions

You can remove the LSM volumes for the system partitions and return to using physical disk partitions. This process is called unencapsulation and involves restarting the system.

Note for TruCluster Environments

To unencapsulate the cluster root domain, use the `volunmigrate` command. See `volunmigrate`(8) and *Cluster Administration* for more information.

The unencapsulation process changes the following files:

- If the root file system is UFS, the `/etc/fstab` file is changed to use disk partitions instead of LSM volumes.
- If the root file system is AdvFS, the `/etc/fdmns/*` directory is updated to change domain directories that have disk partitions associated with the boot disk.
- The `/etc/sysconfigtab` file is changed to update the `swapdevice` entry to not use LSM volumes and to set the `lsm_rootdev_is_volume` entry to 0.

If the system volumes are mirrored, remove all but one plex. Leave only the plex that is using the disk you want the system partitions to use after the unencapsulation completes.

For example, you can remove plexes named `rootvol-01` and `rootvol-02`, leaving the `rootvol` volume with only plex `rootvol-03`, if that plex is on the disk you want to unencapsulate.

The remaining plexes in each system volume can be on different disks from each other; for example, the remaining `rootvol` plex can be on dsk0 while the remaining `swapvol` plex can be on dsk2.

To unencapsulate the system partitions:

1. If the system volumes (root, swap, `/usr` and `/var`) are mirrored, do the following. If not, continue with step 2.

    a. Enter the following command to display volume information:

    ```
    # volprint -v
    ```

    In the output, note the names of the plexes that you want to remove.

    b. Remove all mirror plexes but the one on the disk that you want the system partitions to use after the unencapsulation process completes:

    ```
    # volplex -o rm dis plex-nn
    ```

    For example, to remove secondary plexes for the volumes `rootvol`, `swapvol` and `vol-dsk0g`, enter:

    ```
    # volplex -o rm dis rootvol-02
    # volplex -o rm dis swapvol-02
    # volplex -o rm dis vol-dsk0g-02
    ```

2. Change the boot disk environment variable to point to the physical boot disk (the disk containing the remaining plex for `rootvol`) instead of the LSM volume:

    ```
    # consvar -s bootdef_dev boot_disk
    ```

3. Enter the following command to complete the unencapsulation. This command also removes the LSM private region from the system disks and prompts you to restart the system.

```
# volunroot -a -A
```

Information similar to the following is displayed. Enter **now** at the prompt.

```
This operation will convert the following file systems on the
system/swap disk dsk0 from LSM volumes to regular disk partitions:

        Replace volume rootvol with dsk0a.
        Replace volume swapvol with dsk0b.
        Replace volume vol-dsk0g with dsk0g.
        Remove configuration database on dsk0h.

This operation will require a system reboot.  If you choose to
continue with this operation, your system files will be updated
to discontinue the use of the above listed LSM volumes.
/sbin/volreconfig should be present in /etc/inittab to remove
the named volumes during system reboot.


Would you like to either quit and defer volunroot until later
or commence system shutdown now? Enter either 'quit' or time to be
used with the shutdown(8) command (e.g., quit, now, 1, 5): [quit] now
```

5.4.6.3 Cleaning Up the LSM Disks for Reuse

The disks that were used by the system volumes remain under LSM control as members of the rootdg disk group.

To reuse these disks within LSM or for other purposes, you must remove them from the rootdg disk group and remove the LSM partitions. Then you can remove them from LSM control and either reinitialize them as LSM disks (as sliced disks) or use them for purposes other than LSM.

To clean up the system-specific LSM disks:

1. Display the LSM disks:

```
# voldisk list
```

Information similar to the following is displayed:

```
DEVICE       TYPE      DISK         GROUP        STATUS
dsk0         sliced    -            -            unknown
dsk1a        nopriv    root02       rootdg       online
dsk1b        nopriv    swap02       rootdg       online
dsk1g        nopriv    dsk1g-AdvFS  rootdg       online
dsk1h        simple    dsk1h        rootdg       online
dsk2         sliced    dsk2         rootdg       online
dsk3         sliced    dsk3         rootdg       online
dsk4         sliced    dsk4         rootdg       online
dsk5         sliced    dsk5         rootdg       online
```

2. Remove the disks from the rootdg disk group using their disk media names (in the DISK column):

```
# voldg rmdisk root02 swap02 dsk1g-AdvFS dsk1h
```

3. Remove the disks from LSM control using their disk access names (in the DEVICE column):

```
# voldisk rm dsk1a dsk1b dsk1g dsk1h
```

See Section 4.1.2 for more information on reinitializing these disks for LSM.

5.4.7 Recovering an LSM Volume

You might need to recover an LSM volume that has become disabled. Alert icons and the Alert Monitor window might provide information when an LSM volume recovery is needed. (See the *System Administration* manual for more information about the Alert Monitor.) Recovering an LSM volume starts the disabled volume and, if applicable, resynchronizes mirror plexes or RAID 5 parity.

To recover an LSM volume, enter the following command, specifying either the volume or a disk, if the disk is used by several volumes:

```
# volrecover [-g disk_group] -sb volume|disk
```

The -s option starts all disabled volumes, and the -b option runs the command in the background.

For example, to recover an LSM volume called vol01, enter:

```
# volrecover -sb vol01
```

To recover all LSM objects (subdisks, plexes, or volumes) that use a disk called dsk5, enter:

```
# volrecover -sb dsk5
```

If you do not specify a disk group, LSM volume name, or disk name, all volumes are recovered. If recovery of an LSM volume is not possible, restore the LSM volume from backup. See Section 5.4.3 for more information.

5.4.8 Renaming an LSM Volume

You can rename an LSM volume. The new LSM volume name must be unique within the disk group. If the LSM volume has a file system or is part of an AdvFS domain, you must also update the /etc/fstab and /etc/fdmn files.

To rename an LSM volume, enter:

```
# voledit rename old_volume new_volume
```

Note

If you do not update the relevant files in the `/etc` directory before the system is restarted, subsequent commands using a volume's previous name will fail.

5.4.9 Resizing an LSM Volume

You can increase or decrease the size of an LSM volume; for example, you can increase the size of the primary swap space volume. In LSM, increasing the size of a volume is called *growing a volume* and decreasing its size is called *shrinking a volume*.

Caution

You must be sure that the application using the LSM volume can respond appropriately to growing, or especially shrinking, LSM volumes. You might have to perform additional steps specific to the application using the volume, either before or after changing its size. Refer to the documentation for the application using the volume for more information.

5.4.9.1 Growing a Volume

You can increase the size of a volume by specifying either an amount to grow by or a size to grow to. The size of any log plexes remains unchanged.

Notes on File Systems

If the volume is used for an AdvFS file system, do not increase the space in the domain by growing an underlying LSM volume. Instead, create a new LSM volume and add it to the domain. See *AdvFS Administration* for more information on increasing the size of a domain.

If the volume is used for a file system other than AdvFS, you must perform additional steps specific to the file system type for the file system to take advantage of increased space. See *System Administration* for more information on increasing a file system other than AdvFS.

If an application other than a file system uses the volume, you must make any necessary application modifications after the grow operation is complete.

When growing a volume, you must use the `-f` option to force the change. You can use the `-b` option to perform the operation in the background. This is helpful if the *growto* or *growby* length specified is substantial and if the volume uses mirror plexes or RAID 5, because it will undergo resynchronization as a result of the grow operation.

To grow a volume:

- By a specific amount, enter:

  ```
  # volassist [-g disk_group] -f [-b] growby volume length_change
  ```

- To a new size, enter:

  ```
  # volassist [-g disk_group] -f [-b] growto volume new_length
  ```

5.4.9.2 Shrinking a Volume

You can decrease the size of a volume by specifying either an amount to shrink by, or a size to shrink to. The size of any log plexes remains unchanged.

Cautions

If the volume is used for an AdvFS file system, do not decrease the space in the domain by shrinking an underlying LSM volume. Instead, remove a volume from the domain (in AdvFS, a volume can be a disk, disk partition, or an LSM volume). See *AdvFS Administration* for more information on removing volumes from a domain.

If the volume is used for a file system other than AdvFS, you must perform additional steps specific to the file system type before shrinking the volume, so that the file system can recognize and safely adjust to the decreased space.

There is no direct way to shrink a UFS file system other than backing up the data, destroying the original file system, creating a new file system of the smaller size, and restoring the data into the new file system.

See *System Administration* for more information.

If an application other than a file system uses the volume, you must make any necessary application modifications before shrinking the LSM volume.

When shrinking a volume, you must use the `-f` option to force the change.

To shrink a volume:

- By a specific amount, enter:

```
# volassist [-g disk_group] -f [-b] shrinkby volume length_change
```

- To a new size, enter:

```
# volassist [-g disk_group] -f [-b] shrinkto volume new_length
```

5.4.10 Changing LSM Volume Permission, User, and Group Attributes

By default, the device special files for LSM volumes are created with read and write permissions granted only to the owner. Databases or other applications that perform raw I/O might require device special files to have other settings for the permission, user, and group attributes.

You must use LSM commands to change the permission, user, and group attributes for LSM volumes. The LSM commands ensure that settings for these attributes are stored in the LSM database, which keeps track of all settings for LSM objects.

Do not change the permission, user, or group attributes by using the `chmod`, `chown`, or `chgrp` commands directly on the device special files associated with LSM volumes. These standard UNIX commands do not store the required values in the LSM configuration database.

To change Tru64 UNIX user, group, and permission attributes, enter:

```
# voledit [-g disk_group] set \
user=username group=groupname mode=permission volume
```

The following example changes the user, group, and permission attributes for an LSM volume called vol1:

```
# voledit set user=new_user group=admin mode=0600 vol1
```

5.5 Managing Plexes

The following sections describe how to use LSM commands to manage plexes.

5.5.1 Displaying Plex Information

You can display information about all plexes or about one specific plex.

5.5.1.1 Displaying General Plex Information

To display general information for all plexes, enter:

```
# volprint -pt
```

Information similar to the following is displayed:

```
Disk group: rootdg

PL NAME         VOLUME        KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID MODE
```

```
pl tst-01       tst          ENABLED  ACTIVE   262144   CONCAT    -        RW
pl tst-02       tst          DETACHED STALE    262144   CONCAT    -        RW
pl vol5-01      vol5         ENABLED  ACTIVE   409696   RAID      8/32     RW
pl vol5-02      vol5         ENABLED  LOG      2560     CONCAT    -        RW
```

5.5.1.2 Displaying Detailed Plex Information

To display detailed information about a specific plex, enter:

```
# volprint -lp plex
```

Information similar to the following is displayed:

```
Disk group: rootdg

Plex:   p1
info:   len=500
type:   layout=CONCAT
state:  state=EMPTY kernel=DISABLED io=read-write
assoc:  vol=v1 sd=dsk4-01
flags:  complete

Plex:   p2
info:   len=1000
type:   layout=CONCAT
state:  state=EMPTY kernel=DISABLED io=read-write
assoc:  vol=v2 sd=dsk4-02
flags:  complete

Plex:   vol_mir-01
info:   len=256
type:   layout=CONCAT
state:  state=ACTIVE kernel=ENABLED io=read-write
assoc:  vol=vol_mir sd=dsk2-01
flags:  complete

Plex:   vol_mir-02
info:   len=256
type:   layout=CONCAT
state:  state=ACTIVE kernel=ENABLED io=read-write
assoc:  vol=vol_mir sd=dsk3-01
flags:  complete

Plex:   vol_mir-03
info:   len=0 (sparse)
type:   layout=CONCAT
state:  state=ACTIVE kernel=ENABLED io=read-write
assoc:  vol=vol_mir sd=(none)
flags:
logging: logsd=dsk3-02 (enabled)
```

5.5.2 Adding a Data Plex

You can add a data plex to a volume to create a mirror data plex. You cannot create a mirror data plex on a disk that already contains a data plex for the volume.

The data from the original plex is copied to the added plex, and the plexes are synchronized. This process can take a long time depending on the size of the volume, so you should run the command in the background (using the `&` operator).

Note

Adding a data plex does not add a DRL plex to the volume. It is highly recommended that volumes with mirror plexes have a DRL plex. See Section 5.5.3 for more information on adding a log plex to a volume.

To add a data plex, enter:

```
# volassist mirror volume [disk] &
```

5.5.3 Adding a Log Plex

You can add a log plex (DRL plex or RAID 5 log plex) to a volume that has mirrored data plexes or a RAID 5 data plex. However, if the volume is used for secondary swap, it should not have a DRL. You use the same command to add both DRL plexes and RAID 5 logs.

To improve performance, the DRL plex should not be on the same disk as one of the volume's data plexes. To ensure that LSM does not create a DRL plex on the same disk as a data plex, use the `volprint -ht` command to display volume information to identify an LSM disk that is not part of the volume.

To add a log plex to a volume, enter:

```
# volassist addlog volume [disk]
```

5.5.4 Moving Data to a New Plex

You can move the data from a striped or concatenated plex to a new plex to:

- Move the LSM volume onto disks with better performance.
- Move the LSM volume onto new plexes that use a different data layout type. For example, you can move data from a concatenated plex to a striped plex to improve performance.
- Move the LSM volume to different plexes temporarily so you can repair or replace disks in the original plex.

  You can perform this operation only on volumes that use concatenated or striped plexes. You cannot move data from a concatenated or striped plex to a RAID 5 plex or from a RAID 5 plex to a concatenated or striped plex.

For a move operation to be successful:

- The old plex must be an active part of an active volume.

- The new plex cannot be associated with another LSM volume and must be at least the same size as or larger than the original plex.

Note

If the new plex is larger than the original plex and the original plex contains a file system, the file system will not recognize and use the extra space after it is moved. You must recreate the file system on the new plex to take advantage of the extra space.

To move data from one plex to another:

1. Display the size of the plex you want to move:

   ```
   # volprint -ht volume
   ```

   Information similar to the following is displayed:

   ```
   Disk group: rootdg

   V  NAME         USETYPE      KSTATE   STATE    LENGTH   READPOL   PREFPLEX
   PL NAME         VOLUME       KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID MODE
   SD NAME         PLEX         DISK     DISKOFFS LENGTH   [COL/]OFF DEVICE   MODE

   v  DataVol      fsgen        ENABLED  ACTIVE   204800   SELECT    -
   pl DataVol-01   DataVol      ENABLED  ACTIVE   204800   STRIPE    8/128    RW
   sd dsk0-01      DataVol-01   dsk0     0        25600    0/0       dsk0     ENA
   sd dsk1-01      DataVol-01   dsk1     0        25600    1/0       dsk1     ENA
   sd dsk2-01      DataVol-01   dsk2     0        25600    2/0       dsk2     ENA
   sd dsk3-01      DataVol-01   dsk3     0        25600    3/0       dsk3     ENA
   sd dsk4-01      DataVol-01   dsk4     0        25600    4/0       dsk4     ENA
   sd dsk6-01      DataVol-01   dsk6     0        25600    5/0       dsk6     ENA
   sd dsk7-01      DataVol-01   dsk7     0        25600    6/0       dsk7     ENA
   sd dsk8-01      DataVol-01   dsk8     0        25600    7/0       dsk8     ENA
   pl DataVol-02   DataVol      ENABLED  ACTIVE   204800   STRIPE    8/128    RW
   sd dsk10-01     DataVol-02   dsk10    0        25600    0/0       dsk10    ENA
   sd dsk11-01     DataVol-02   dsk11    0        25600    1/0       dsk11    ENA
   sd dsk12-01     DataVol-02   dsk12    0        25600    2/0       dsk12    ENA
   sd dsk13-01     DataVol-02   dsk13    0        25600    3/0       dsk13    ENA
   sd dsk14-01     DataVol-02   dsk14    0        25600    4/0       dsk14    ENA
   sd dsk15-01     DataVol-02   dsk15    0        25600    5/0       dsk15    ENA
   sd dsk18-01     DataVol-02   dsk18    0        25600    6/0       dsk18    ENA
   sd dsk19-01     DataVol-02   dsk19    0        25600    7/0       dsk19    ENA
   pl DataVol-03   DataVol      ENABLED  ACTIVE   LOGONLY  CONCAT    -        RW
   sd dsk20-02     DataVol-03   dsk20    0        65       LOG       dsk20    ENA
   ```

   In this example, the volume has two striped data plexes of 204800 sectors (100 MB).

2. Ensure there is enough space on other LSM disks to move the plex's data.

3. Create a new plex with the characteristics you want.

   - For a concatenated plex, see Section 4.2.1
   - For a striped plex, see Section 4.2.2

- For a striped plex that uses disks on different buses, see Section 4.2.3

4. Enter the following command line (set to run in the background) to attach the new plex to the volume and move the data from the old plex to the new plex, optionally removing the old plex upon successful completion of the move:

   ```
   # volplex [-o rm] mv old_plex new_plex &
   ```

   The volume remains active and usable during this operation.

5.5.5 Reattaching a Plex

If you removed a plex from a volume and did not recursively remove it and its objects, you can reattach it to the volume.

To reattach a plex to a volume, enter:

```
# volplex att volume plex
```

5.5.6 Removing a Plex

You can remove a plex from an LSM volume to reduce the number of plexes in a volume.

Note

The following restrictions apply:

- You cannot remove a RAID 5 data plex from a volume, because that is the only data plex. However, you can remove a volume completely (Section 5.4.6).
- If the volume has mirror plexes and you remove all but one, the volume's data is no longer redundant.
- If you remove the DRL plex from a volume that has mirror plexes and the system fails, LSM will have to resynchronize the entire contents of the plexes when the system restarts.
- If you remove the RAID 5 log plex from a volume that uses a RAID 5 plex and the system fails, LSM will have to read back all the volume data, regenerate the parity for each stripe, and rewrite each stripe in the plex.

To remove a data plex from a volume with mirror plexes:

1. Dissociate the plex from its volume, and optionally remove the old plex after successful completion of the dissociation:

   ```
   # volplex [-o rm] dis plex
   ```

2. If you did not use the option in step 1, remove the plex:

   ```
   # voledit -r rm plex
   ```

Removing the plex also removes all associated subdisks in that plex. The disks remain under LSM control, and you can use them for other volumes or remove them from LSM control.

To remove the log plex from a RAID 5 volume:

1. Dissociate the log plex from the RAID 5 volume (using the `-f` option):

   ```
   # volplex -f dis log_plex
   ```

2. Remove the plex and its subdisks:

   ```
   # voledit -r rm log_plex
   ```

5.6 Managing Subdisks

The following sections describe how to use LSM commands to manage subdisks.

5.6.1 Displaying Subdisk Information

You can display information about all subdisks or one specific subdisk.

5.6.1.1 Displaying General Subdisk Information

To display general information for all subdisks, enter:

```
# volprint -st
```

Information similar to the following is displayed:

```
Disk group: rootdg

SD NAME      PLEX        DISK   DISKOFFS    LENGTH  [COL/]OFF DEVICE  MODE

sd dsk2-01   vol_mir-01  dsk2   0            256        0      dsk2    ENA
sd dsk3-02   vol_mir-03  dsk3   0             65      LOG      dsk3    ENA
sd dsk3-01   vol_mir-02  dsk3   65           256        0      dsk3    ENA
sd dsk4-01   p1          dsk4   17           500        0      dsk4    ENA
sd dsk4-02   p2          dsk4   518         1000        0      dsk4    ENA
```

5.6.1.2 Displaying Detailed Subdisk Information

To display detailed information about a specific subdisk, enter:

```
# volprint -l subdisk
```

The following example shows information about a subdisk called dsk12–01:

```
Disk group: rootdg

Subdisk:  dsk12-01
```

```
info:     disk=dsk12 offset=0 len=2560
assoc:    vol=vol5 plex=vol5-02 (offset=0)
flags:    enabled
device:   device=dsk12 path=/dev/disk/dsk12g diskdev=82/838
```

5.6.2 Joining Subdisks

You can join two or more subdisks to form a single, larger subdisk. Subdisks can be joined only if they belong to the same plex and occupy adjacent regions of the same disk. For a volume with striped plexes, the subdisks must be in the same column. The joined subdisk can have a new subdisk name or retain the name of one of the subdisks being joined.

To join subdisks, enter:

```
# volsd join subdisk1 subdisk2 new_subdisk
```

5.6.3 Splitting Subdisks

You can divide a subdisk into two smaller subdisks. Once split, you can move the data in the smaller subdisks to different disks. This is useful for reorganizing volumes or for improving performance. The new, smaller subdisks occupy adjacent regions within the same region of the disk that the original subdisk occupied.

You must specify a size for the first subdisk; the second subdisk consists of the rest of the space in the original subdisk.

If the subdisk to be split is associated with a plex, both of the resultant subdisks are associated with the same plex. You cannot split a log subdisk.

To split a subdisk and assign each subdisk a new name, enter:

```
# volsd -s size split original_subdisk new_subdisk1 new_subdisk2
```

To split a subdisk and retain the original name for the first subdisk and assign a new name to the second subdisk, enter:

```
# volsd -s size split original_subdisk new_subdisk
```

5.6.4 Moving Subdisks to a Different Disk

You can move the data in subdisks to a different disk to improve performance. The disk space occupied by the data in the original subdisk is returned to the free space pool.

Ensure that the following conditions are met before you move data in a subdisk:

- Both source and destination subdisks must be the same size.
- The source subdisk must be part of an active plex on an active volume.

- The destination subdisk must not be associated with any other plex.

To move data from one subdisk to another, enter:

```
# volsd mv source_subdisk target_subdisk
```

5.6.5 Removing a Subdisk

You can remove a subdisk that is not associated with or needed by an LSM volume. Removing a subdisk returns the disk space to the free space pool in the disk group. To remove a subdisk, you must dissociate the subdisk from a plex, then remove it.

To remove a subdisk:

1. Display information about the subdisk to identify any volume or plex associations:

   ```
   # volprint -l subdisk
   ```

   - If the subdisk is associated with a volume, information similar to the following is displayed:

     ```
     Disk group: rootdg

     Subdisk:  dsk9-01
     info:     disk=dsk9 offset=0 len=2048
     assoc:    vol=newVol plex=myplex (column=1 offset=0)
     flags:    enabled
     device:   device=dsk9 path=/dev/disk/dsk9g diskdev=82/646
     ```

   - If the subdisk has no associations to any plex or volume, information similar to the following is displayed:

     ```
     Disk group: dg1

     Subdisk:  dsk5-01
     info:     disk=dsk5 offset=0 len=2046748
     assoc:    vol=(dissoc) plex=(dissoc)
     flags:    enabled
     device:   device=dsk5 path=/dev/disk/dsk5g diskdev=79/390
     ```

2. Do one of the following to remove the subdisk:

   - If the subdisk is associated with a volume, enter:

     ```
     # volsd -o rm dis subdisk
     ```

   - If the subdisk is not part of a volume and has no associations, enter:

     ```
     # voledit [-g disk_group] rm subdisk
     ```

6 Troubleshooting

LSM helps you protect the availability and reliability of data but does not prevent I/O failure. LSM is simply another layer added to the I/O subsystem. LSM depends on the underlying disk device drivers and system files to decide on the availability of individual disks and to manage and report any failures.

This chapter describes how to troubleshoot common LSM problems, describes tools that you can use to learn about problems, and offers possible solutions.

The hot-spare feature provides the best protection for volumes that use mirror plexes or a RAID 5 plex. When enabled, the hot-spare feature allows LSM to automatically relocate data from a failed disk in a volume that uses either a RAID 5 plex or mirrored plexes. LSM writes the data to a designated spare disk, or to free disk space, and sends you mail about the relocation. See Section 3.4.4.1 for more information about enabling the hot-spare feature.

6.1 Monitoring LSM

You can use LSM commands to monitor the status of LSM objects. By doing so, you can understand how LSM works under normal conditions and watch for indication that an LSM object might need adjustments before a problem arises.

6.1.1 Monitoring LSM Events

By default, LSM uses Event Manager (EVM) software to log events. The events that LSM logs are defined in the EVM template called `/usr/share/evm/templates/sys/lsm.volnotify.evt`.

You can select, filter, sort, format, and display LSM events using EVM commands or the graphical event viewer, which is integrated with the SysMan Menu and SysMan Station.

To display a list of logged LSM events, enter:

```
# evmget -f "[name *.volnotify]" | evmshow -t "@timestamp @@"
```

To display LSM events in real time, enter:

```
# evmwatch -f "[name *.volnotify]" | evmshow -t "@timestamp @@"
```

See the EVM(5) reference page for more information about EVM.

6.1.2 Monitoring Read and Write Statistics

You can use the volstat command to view and reset:

- The number of successful or failed read and write operations
- The number of blocks read from and written to
- The average time spent on read and write operations. This time reflects the total time it took to complete a read or write operation, including the time spent waiting in a queue on a busy device.

Table 6–1 describes the some of the options that you can use with the volstat command.

Table 6–1: Common volstat Command Options

Option	Displays
-v	Volume statistics
-p	Plex statistics
-s	Subdisk statistics
-d	LSM disk statistics
-i *seconds*	The specified statistics continuously in the interval specified (in seconds).

For information on all the volstat options, see the volstat(8) reference page.

Note

In a cluster environment, the volstat command displays statistics for the system on which the command is entered and does not provide statistics for all the systems within a cluster.

6.1.2.1 Displaying Read and Write Statistics

To display read and write statistics for LSM objects, enter:

```
# volstat [-g disk_group] -vpsd [-i number_of_seconds]
```

Information similar to the following is displayed:

```
                OPERATIONS          BLOCKS        AVG TIME(ms)
TYP NAME      READ     WRITE      READ     WRITE   READ  WRITE
dm  dsk6        3        82        40     62561    8.9   51.2
dm  dsk7        0       725         0    176464    0.0   16.3
dm  dsk9      688        37    175872       592    3.9    9.2
dm  dsk10   29962         0   7670016         0    4.0    0.0
dm  dsk12       0     29962         0   7670016    0.0   17.8
vol v1          3        72        40     62541    8.9   56.5
pl  v1-01       3        72        40     62541    8.9   56.5
sd  dsk6-01     3        72        40     62541    8.9   56.5
vol v2          0        37         0       592    0.0   10.5
pl  v2-01       0        37         0       592    0.0    8.0
sd  dsk7-01     0        37         0       592    0.0    8.0
sd  dsk12-01    0         0         0         0    0.0    0.0
pl  v2-02       0        37         0       592    0.0    9.2
sd  dsk9-01     0        37         0       592    0.0    9.2
sd  dsk10-01    0         0         0         0    0.0    0.0
pl  v2-03       0         6         0        12    0.0   13.3
sd  dsk6-02     0         6         0        12    0.0   13.3
```

The LSM objects are identified as follows:

- `dm` – Disk media name (LSM name for the disk)
- `vol` – Volume name
- `pl` – Plex name
- `sd` – Subdisk name

6.1.2.2 Displaying Failed Read and Write Statistics

To display failed I/O statistics, enter:

`#` **`volstat`** [**`-g`** ***`disk_group`***] **`-f cf`** ***`LSM_object`***

Information similar to the following is displayed:

```
                CORRECTED          FAILED
TYP NAME        READS   WRITES    READS   WRITES
vol testvol         1        0        0        0
```

LSM corrects read failures for mirror plexes or a RAID 5 plex, because these plexes provide data redundancy.

6.1.3 Monitoring LSM Object States

The kernel and LSM monitor the state of LSM objects.

To display the state of LSM objects, enter:

`#` **`volprint`** [**`-g`** ***`disk_group`***]

Information similar to the following is displayed:

```
.
.
.
Disk group: dg1
```

```
TY NAME         ASSOC        KSTATE   LENGTH   PLOFFS   STATE    TUTIL0  PUTIL0
dg dg1          dg1          -        -        -        -        -       -

dm dsk1         dsk1         -        2046748  -        -        -       -
dm dsk2         dsk2         -        2046748  -        -        -       -
dm dsk4         dsk4         -        2046748  -        -        -       -
dm dsk5         dsk5         -        2046748  -        -        -       -

v  vol-test     fsgen        ENABLED  2048     -        ACTIVE   -       -
pl vol-test-01  vol-test     ENABLED  2048     -        ACTIVE   -       -
sd dsk1-01      vol-test-01  ENABLED  1024     0        -        -       -
sd dsk2-01      vol-test-01  ENABLED  1024     0        -        -       -
pl vol-test-02  vol-test     ENABLED  2048     -        ACTIVE   -       -
sd dsk4-01      vol-test-02  ENABLED  1024     0        -        -       -
sd dsk5-01      vol-test-02  ENABLED  1024     0        -        -       -
```

The KSTATE column shows the kernel state of the LSM object. The STATE column shows the LSM state of the LSM object. The LSM objects are identified as follows:

- `dg` – Disk group name
- `dm` – Disk media name (LSM name for the disk)
- `v` – Volume name
- `pl` – Plex name
- `sd` – Subdisk name

6.1.3.1 LSM Kernel States

The LSM kernel state indicates the accessibility of the LSM object as viewed by the kernel. Table 6–2 describes kernel states for LSM objects.

Table 6–2: LSM Volume Kernel States (KSTATE)

Kernel State	Means
ENABLED	The LSM object is accessible and read and write operations can be performed.
DISABLED	The LSM object is not accessible.
DETACHED	Read and write operations cannot be performed, but device operations are accepted.

6.1.3.2 LSM Object States

LSM monitors the states of volumes, plexes, and subdisks.

- A volume has an LSM state (Table 6–3). The meaning of some volume states differs depending on the kernel state (KSTATE).
- A plex has an LSM state (Table 6–4).
- A subdisk has an LSM state (Table 6–5).

Table 6–3: LSM Volume States (STATE)

State	Means	Kernel State
EMPTY	The volume contents are not initialized.	DISABLED
CLEAN	The volume is not started. • For mirrored volumes, plexes are synchronized. • For RAID 5 volumes, parity is good and stripes are consistent.	DISABLED
ACTIVE	The volume was started or was in use when the system was restarted.	ENABLED DISABLED if RAID 5 parity synchronization is not guaranteed or if mirror plexes are not guaranteed to be consistent.
SYNC	The system is resynchronizing mirror plexes or RAID 5 parity.	ENABLED if the mirror plexes or RAID 5 parity are being resynchronized. DISABLED if the mirror plexes or RAID 5 parity were being resynchronized when the system restarted and therefore still need to be synchronized.
NEEDSYNC	The volume requires a resynchronization operation the next time it starts.	
REPLAY	A RAID 5 volume is in a transient state as part of a log replay. A log replay occurs when it is necessary to reconstruct data using parity and data.	

Table 6–4: LSM Plex States

State	Means
EMPTY	The plex is not initialized. This state is also set when the volume state is EMPTY.
CLEAN	The plex was running normally when the volume was stopped. The plex was enabled without requiring recovery when the volume was started.
ACTIVE	The plex is running normally on a started volume.
LOG	The plex is a DRL or RAID 5 log plex for the volume.

Table 6–4: LSM Plex States (cont.)

State	Means
STALE	The plex was detached, either by the `volplex det` command or by an I/O failure. STALE plexes are reattached automatically by `volplex att` when a volume starts.
IOFAIL	The `vold` daemon places an ACTIVE plex in the IOFAIL state when it detects an error. The plex is disqualified from the recovery selection process at volume start time, ensuring that LSM uses only valid plexes for recovery. A plex marked IOFAIL is recovered if possible during a resynchronization.
OFFLINE	The plex was disabled by the `volmend off` command.
SNAPATT	This is a snapshot plex that is attached by the `volassist snapstart` command. When the attach is complete, the state for the plex is changed to SNAPDONE. If the system fails before the attach completes, the plex and all of its subdisks are removed.
SNAPDONE	This is a snapshot plex created by `volassist snapstart` command that is fully attached. You can turn a plex in this state into a snapshot volume with the `volassist snapshot` command. If the system fails before the attach completes, the plex and all of its subdisks are removed.
SNAPTMP	This is a snapshot plex that is attached by the `volplex snapstart` command. When the attach is complete, the state for the plex changes to SNAPDIS. If the system fails before the attach completes, the plex is dissociated from the volume.
SNAPDIS	This is a snapshot plex created by the `volplex snapstart` command that is fully attached. You can turn a plex in this state into a snapshot volume with the `volplex snapshot` command. If the system fails before the attach completes, the plex is dissociated from the volume.
TEMP	This is a plex that is associated and attached to a volume with the `volplex att` command. If the system fails before the attach completes, the plex is dissociated from the volume.
TEMPRM	This is a plex that is being associated and attached to a volume with the `volplex att` command. If the system fails before the attach completes, the plex is dissociated from the volume and removed. Any subdisks in the plex are kept.
TEMPRMSD	This is a plex that is being associated and attached to a volume with the `volplex att` command. If the system fails before the attach completes, the plex and its subdisks are dissociated from the volume and removed.

Table 6–5: LSM Subdisk States

State	Means
REMOVED	The subdisk (which might encompass the entire LSM disk) was removed from the volume, disk group, or from LSM control.
RECOVER	The subdisk must be recovered. Use the `volrecover` command.

6.2 Missing or Altered sysconfigtab File

During the boot disk encapsulation procedure, LSM adds the following entries to the `/etc/sysconfigtab` file to enable the system to boot off the LSM root volume:

```
lsm:
lsm_rootdev_is_volume=1
```

If this file is deleted or the LSM-specific entries are deleted, the system will not boot. If this happens, do the following:

1. Boot the system interactively:

   ```
   >>> boot -fl i
   .........
   .........
   Enter kernel_name option_1 ... option_n: vmunix
   ```

2. Reset the LSM entries as follows:

   ```
   lsm:
   lsm_rootdev_is_volume=1
   ```

6.3 LSM Startup and Command Problems

LSM requires that the `vold` and `voliod` daemons be running. These daemons are normally started automatically when the system boots. If these daemons are not running, the most obvious problems you might notice are that LSM commands fail to complete or do not respond as expected, which is an indication that LSM did not correctly start up.

The following sections describe how to check if the daemons are running and how to correct problems.

6.3.1 Checking the vold Daemon

To determine the state of the `vold` daemon, enter:

```
# voldctl mode
```

Table 6–6 shows messages that might display, what the message means, and the commands you should enter if `vold` is disabled or not running.

Table 6–6: vold Messages and Solutions

Message	Status	Enter
`Mode: enabled`	Running and enabled	—
`Mode: disabled`	Running but disabled	**`voldctl enable`**
`Mode: not-running`	Not running	**`vold`**

See the `vold`(8) reference page for more information on the `vold` daemon.

6.3.2 Checking the voliod Daemon

The correct number of `voliod` daemons automatically start when LSM starts. Typically several `voliod` daemons are running at all times. You should run at least one `voliod` daemon for each processor on the system.

To display the number of the `voliod` daemons running, enter:

```
# voliod
```

Information similar to the following is displayed:

```
2 volume I/O daemons running
```

This is the only method for checking `voliod` daemons, because the `voliod` processes are kernel threads and do not display in the output of the `ps` command.

If no `voliod` daemons are running, or if you want to change the number of daemons, enter the following command where *n* is the number of I/O daemons to start:

```
# voliod set n
```

Set the number of LSM I/O daemons to two or the number of central processing units (CPUs) on the system, whichever is greater.

See the `voliod`(8) reference page for more information on the `voliod` daemon.

6.4 Solving Problems with LSM Volumes

The following sections describe how to solve common LSM volume problems.

6.4.1 Insufficient Space to Create a Volume

When you use the `volassist` command to create a volume with a striped plex, you might receive an error message indicating insufficient space for the volume even though you know there is enough space available.

The `volassist` command rounds up the length you specify on the command line to a multiple of the data unit size of 64 KB by default, or the stripe width you specified, and then divides the total by the number of disks available to make the column. The smallest disk in the disk group limits the data unit size.

For example, you have two disks with differing free space in the disk group called dg1:

```
# voldg -g dg1 free
```

```
GROUP  DISK      DEVICE    TAG      OFFSET   LENGTH    FLAGS
dg1    dsk1      dsk1      dsk1     0        2049820  -
dg1    dsk2      dsk2      dsk2     0        2047772  -
```

The total free space on these two disks is 4097592. You tried to create a volume with a striped plex, with a length of 4095544, or about 2 GB, which is less than the total space available:

```
# volassist -g dg1 make NewVol 4095544 layout=stripe
```

```
volassist: adjusting length 4095544 to conform
to a layout of    2    stripes    128    blocks wide
volassist: adjusting length up to 4095744 blks

volassist: insufficient space for a 4095744 block long volume in stripe,
contiguous layout
```

The command returned an error message indicating insufficient space, because `volassist` rounded up the length you specified to an even multiple of the data unit size of 64 KB (128 blocks) and divided that number by the number of disks (2). The result was larger than the space available on the smaller disk: 4095744 / 2 = 2048796.

If your volume does not need to be precisely the size you specified, you can retry the command with a length that works with the data unit size and the number of disks. For example, multiply the size of the smallest free space by the number of disks: 2047772 * 2 = 4095488. Use this value in the command line:

```
# volassist -g dg1 make NewVol 4095488 layout=stripe
```

If the volume you require is larger than the total free space in the disk group, or if the volume must be exactly the size you specify, you must add more (or larger) disks to that disk group. See Section 5.2.3 for more information on adding disks to a disk group.

6.4.2 Starting a Disabled Volume

If you cannot mount a file system or open an LSM volume, the LSM volume might not be started.

To determine whether or not the LSM volume is started, enter:

volinfo [**-g** ***disk_group***] ***volume***

The following output shows the condition of several volumes:

```
vol   bigvol         fsgen    Startable
vol   vol2           fsgen    Started
vol   brokenvol      gen      Unstartable
```

LSM volumes can have the following conditions:

- `Started` – The volume is enabled and running normally.
- `Startable` – The volume is not enabled, and at least one plex has a state of ACTIVE or CLEAN, indicating that the volume can be restarted.[1]

  To start a startable volume, enter:

  # **volume** [**-g** ***disk_group***] **start** ***volume***

- `Unstartable` – The volume is not enabled and has a problem that you must resolve before you can start the volume. For example, a disk might have failed.
  - If the volume is redundant (that is, it uses mirror plexes or a RAID 5 plex), see Section 6.5.5 for information on replacing failed disks and recovering the volumes.
  - If the volume is not redundant, see Section 6.4.3.

6.4.3 Recovering Unstartable Nonredundant Volumes

Nonredundant volumes are those that use a single plex that is either concatenated or striped. If a disk in the plex fails, the volume will be unstartable.

You can display the volume's condition by entering:

volinfo -p

Information similar to the following is displayed:

```
vol  tst           fsgen    Unstartable
plex tst-01        NODEVICE
```

To recover the volume:

1. If the disk is usable, continue with step 2. If the disk has failed, replace the disk:

   a. Identify the disk media name of the failed disk using one of the following commands:

---

[1] Normally, volumes will not be in this state unless you manually created a volume (not using the `volassist` command, which starts the new volume automatically), or you did something that disabled the volume, such as removing a plex. All startable volumes are started when the system restarts.

- To display all disk, disk group, and volume information and the status of any volumes that are affected by the failed disk, enter:

  ```
  # volprint -Aht
  ```

- To display only the disk information, enter:

  ```
  # volprint -Adt
  ```

b. Remove the failed disk and retain the disk media records:

   ```
   # voldg [-g disk_group] -k rmdisk disk_media_name
   ```

c. Remove the disk from LSM control, using the disk access name:

   ```
   # voldisk rm disk_access_name
   ```

d. Physically remove the failed disk and replace it with the new disk.

   Please note that the device must be completely removed from LSM before running any non-LSM commands to remove and replace the failed disk, such as `hwmgr -redirect`.

e. Scan for the new disk:

   ```
   # hwmgr -scan scsi
   ```

   The `hwmgr` command returns the prompt before it completes the scan. You need to know that the system has discovered the new disk before continuing. See the `hwmgr`(8) reference page for more information on how to trap the end of a scan.

f. Label and initialize the new disk:

   - If you have a backup of the previous disk's disk label information (Section 4.1.3):

     i. Apply the backup disk label to the new disk:

        ```
        # disklabel -R disk_access_name auto file
        ```

     ii. Initialize the disk for LSM, using the disk access name:

        ```
        # voldisk -f init disk_access_name
        ```

   - If no disk label file is available:

     i. Apply a default disk label to the new disk:

        ```
        # disklabel -rwn disk_access_name
        ```

     ii. Initialize the new disk for LSM:

        ```
        # voldisksetup -i disk
        ```

g. Optionally (but recommended), create a backup copy of the new disk's disk label information:

   ```
   # disklabel disk_access_name > file
   ```

h. Add the new disk to the applicable disk group, assigning a disk media name to the disk access name. You can reuse the disk media name of the failed disk as the disk media name for the new disk:

```
# voldg [-g disk_group] -k adddisk
disk_media_name=disk_access_name
```

i. Verify that the volume's plex state has changed to RECOVER:

```
# volinfo -p
vol  tst            fsgen    Unstartable
plex tst-01         RECOVER
```

2. Set the plex state to STALE:

```
# volmend fix stale plex
```

LSM has internal state restrictions that require a plex to change states in a specific order. A plex must be STALE before it can be marked CLEAN.

3. Set the plex state to CLEAN:

```
# volmend fix clean plex
```

4. Start the volume:

```
# volume start volume
```

The volume is now running and usable but contains invalid data.

5. Depending on what was using the volume, do one of the following:

   - If the volume was used by a file system, recreate the file system on the volume, and mount the file system. See Section 4.3 for more information on configuring a volume for a file system.

     If you have a backup of the data, restore the volume using the backup. See Section 5.4.3 for more information on restoring a volume from backup.

   - If you have no backup and the volume was used by an application such as a database, refer to that application's documentation for information on restoring or recreating the data.

6.4.4 Recovering Volumes with Mirror Plexes

Volumes with mirror plexes are less vulnerable than volumes with a single (nonmirrored) plex, but if disks in all the plexes fail, the volume's state will be Unstartable.

There are three possible scenarios for the failure and recovery of volumes with mirror plexes:

- Data in all plexes is known to be bad or is unknown. See Section 6.4.4.1.

- One plex is known to be valid, and you want to use that plex to restore the others. See Section 6.4.4.2.
- Data in all the plexes is known to be valid, but you have lost all copies of the configuration database. (All disks containing copies failed.) See Section 6.4.4.3.

6.4.4.1 Recovering a Volume with No Valid Plexes

If disks in multiple plexes of a volume failed, all the volume's data might be corrupt or suspect. Recovering a volume from a multiple disk failure requires that you restore the data from backup.

To recover a volume with no valid plexes:

1. Set all the plexes in the volume to CLEAN:

   ```
   # volmend fix clean plex1 plex2 ...
   ```

2. Start the volume:

   ```
   # volume start volume
   ```

3. Depending on what was using the volume, do one of the following:

   - If the volume was used by a file system, recreate the file system on the volume, and mount the file system. See Section 4.3 for more information on configuring a volume for a file system.

     If you have a backup of the data, restore the volume using the backup. See Section 5.4.3 for more information on restoring a volume from backup.

   - If you have no backup and the volume was used by an application such as a database, refer to that application's documentation for information on restoring or recreating the data.

6.4.4.2 Recovering a Volume with One Valid Plex

If you know that one plex in a volume contains valid data, you can use that plex to restore the others.

To recover a volume with one valid plex:

1. Set the valid plex's state to CLEAN:

   ```
   # volmend fix clean valid_plex
   ```

2. Set the state of all the other plexes to STALE:

   ```
   # volmend fix stale stale_plex1 stale_plex2 ...
   ```

3. Start the volume and initiate the resynchronization process in the background:

```
# volrecover -sb volume
```

6.4.4.3 Recovering a Volume After Loss of the Configuration Database

The following procedure requires a backup copy of the configuration database (created by the `volsave` command, as described in Section 5.3.1) restored by the `volrestore` command. See Section 5.3.2 for more information on restoring the configuration database. You should have a high degree of confidence that the volume data is still valid.

To recover a volume after restoring the configuration database:

1. Set all plexes in the volume to CLEAN:

   ```
   # volmend fix clean plex1 plex2 ...
   ```

2. Start the volume:

   ```
   # volume start volume
   ```

6.4.5 Recovering Volumes with a Failed RAID 5 Plex

Volumes that use a RAID 5 plex are designed to remain available when one disk fails. However, if two disks in the data plex fail, the entire volume is compromised.

- If hot-sparing is enabled at the time of a disk failure, system administrator intervention is not required (unless there is no suitable disk space available for relocation). Hot-sparing is triggered by the disk failure, and you are notified of the failure by electronic mail. Hot-sparing automatically attempts to relocate the subdisks of a failing RAID 5 plex. After relocation takes place, the hot-sparing daemon (`volspared`) also initiates a parity resynchronization.

  In the case of a failing RAID 5 log plex, relocation occurs only if the log plex is mirrored; `volspared` then initiates a mirror resynchronization to recreate the RAID 5 log plex.

- If hot-sparing is disabled at the time of a failure, you might need to initiate a resynchronization or recovery.

There are three possible scenarios for the failure and recovery of volumes with a RAID 5 plex:

- Within the data plex, a disk in one column fails. See Section 6.5.5 for more information on replacing a failed disk and recovering the volume.
- Within the data plex, disks in two or more columns fail. See Section 6.4.5.1.
- Within the log plex, a disk fails, or the log plex becomes detached. See Section 6.4.5.2.

6.4.5.1 Recovering a RAID 5 Plex from Multiple Disk Failures

If disks in two or more columns of a RAID 5 data plex fail, LSM cannot use the remaining data (if any) and parity to reconstruct the missing data. You must restore the data from backup.

To restore the volume:

1. If the disk is usable, continue with step 2. If the disk has failed, replace the disk:

   a. Identify the disk media name of the failed disk using one of the following commands:

      - To display all disk, disk group, and volume information and the status of any volumes that are affected by the failed disk, enter:

        ```
        # volprint -Aht
        ```

      - To display only the disk information, enter:

        ```
        # volprint -Adt
        ```

   b. Remove the failed disk and retain the disk media records:

      ```
      # voldg [-g disk_group] -k rmdisk disk_media_name
      ```

   c. Remove the disk access records, using the disk access name:

      ```
      # voldisk rm disk_access_name
      ```

   d. Physically remove the failed disk and replace it with the new disk.

      Please note that the device must be completely removed from LSM before running any non-LSM commands to remove and replace the failed disk, such as `hwmgr -redirect`.

   e. Scan for the new disk:

      ```
      # hwmgr -scan scsi
      ```

      The `hwmgr` command returns the prompt before it completes the scan. You need to know that the system has discovered the new disk before continuing. See the `hwmgr`(8) reference page for more information on how to trap the end of a scan.

   f. Label and initialize the new disk:

      - If you have a backup of the previous disk's disk label information (Section 4.1.3):

        i. Apply the backup disk label to the new disk:

           ```
           # disklabel -R disk_access_name auto file
           ```

        ii. Initialize the disk for LSM, using the disk access name:

           ```
           # voldisk -f init disk_access_name
           ```

- If no disk label file is available:

  i. Apply a default disk label to the new disk:

  ```
  # disklabel -rwn disk_access_name
  ```

  ii. Initialize the disk for LSM:

  ```
  # voldisksetup -i disk
  ```

g. Optionally (but recommended), create a backup copy of the new disk's disk label information:

```
# disklabel disk_access_name > file
```

h. Add the new disk to the applicable disk group, assigning a disk media name to the disk access name. You can reuse the disk media name of the failed disk as the disk media name for the new disk:

```
# voldg [-g disk_group] -k adddisk
disk_media_name=disk_access_name
```

i. Verify that the volume's plex state has changed to RECOVER:

```
# volinfo -p
vol  tst           fsgen    Unstartable
plex tst-01        RECOVER
```

2. Stop the volume:

```
# volume stop volume
```

3. Set the RAID 5 data plex state to EMPTY:

```
# volmend -f fix empty volume
```

Setting the plex state to EMPTY causes LSM to recalculate the parity when you restart the volume in the next step.

4. Start the volume. The process of recalculating the parity can take a long time; you can run this operation in the background to return the system prompt immediately:

```
# volume [-o bg] start volume
```

The volume becomes usable even while the parity regeneration is underway. If users access a region of the volume that has not yet had its parity recalculated, LSM recalculates the parity for the entire stripe that contains the accessed data before honoring the read or write request.

5. Depending on what was using the volume, do one of the following:

- If the volume was used by a file system, recreate the file system on the volume, and mount the file system. See Section 4.3 for more information on configuring a volume for a file system.

If you have a backup of the data, restore the volume using the backup. See Section 5.4.3 for more information on restoring a volume from backup.

- If you have no backup, and the volume was used by an application such as a database, refer to that application's documentation for information on restoring or recreating the data.

6.4.5.2 Recovering a RAID 5 Log Plex

A disk containing a RAID 5 log could experience a failure. This has no direct effect on the operation of the volume; however, the loss of all RAID 5 logs on a volume makes the volume vulnerable to a complete failure.

The following output from the `volprint` command shows a failure within a RAID 5 log plex. The plex state is BADLOG, and the RAID 5 log plex vol5–02 has failed.

```
Disk group: rootdg

V  NAME         USETYPE      KSTATE   STATE    LENGTH   READPOL   PREFPLEX
PL NAME         VOLUME       KSTATE   STATE    LENGTH   LAYOUT    NCOL/WID MODE
SD NAME         PLEX         DISK     DISKOFFS LENGTH   [COL/]OFF DEVICE   MODE

v  vol5         raid5        ENABLED  ACTIVE   409696   RAID      -
pl vol5-01      vol5         ENABLED  ACTIVE   409696   RAID      8/32     RW
sd dsk3-01      vol5-01      dsk3     0        58528    0/0       dsk3     ENA
sd dsk4-01      vol5-01      dsk4     0        58528    1/0       dsk4     ENA
sd dsk5-01      vol5-01      dsk5     0        58528    2/0       dsk5     ENA
sd dsk6-01      vol5-01      dsk6     0        58528    3/0       dsk6     ENA
sd dsk7-01      vol5-01      dsk7     0        58528    4/0       dsk7     ENA
sd dsk8-01      vol5-01      dsk8     0        58528    5/0       dsk8     ENA
sd dsk9-01      vol5-01      dsk9     0        58528    6/0       dsk9     ENA
sd dsk10-01     vol5-01      dsk10    0        58528    7/0       dsk10    ENA
pl vol5-02      vol5         DISABLED BADLOG   2560     CONCAT    -        RW
sd dsk11-01     vol5-02      dsk11    0        2560     0         -        RMOV
```

RAID 5 log plexes might have a state of DETACHED due to disk failures.

To recover a RAID 5 log plex:

1. If the disk is usable but the log plex is detached, continue with step 2. If the disk has failed, replace the disk:

   a. Identify the disk media name of the failed disk using one of the following commands:

      - To display all disk, disk group, and volume information and the status of any volumes that are affected by the failed disk, enter:

        ```
        # volprint -Aht
        ```

      - To display only the disk information, enter:

        ```
        # volprint -Adt
        ```

   b. Remove the failed disk and retain the disk media records:

voldg [**-g** ***disk_group***] **-k rmdisk** ***disk_media_name***

c. Remove the disk access records, using the disk access name:

   # **voldisk rm** ***disk_access_name***

d. Physically remove the failed disk and replace it with the new disk.

   Please note that the device must be completely removed from LSM before running any non-LSM commands to remove and replace the failed disk, such as `hwmgr -redirect`.

e. Scan for the new disk:

   # **hwmgr -scan scsi**

   The `hwmgr` command returns the prompt before it completes the scan. You need to know that the system has discovered the new disk before continuing. See the `hwmgr`(8) reference page for more information on how to trap the end of a scan.

f. Label and initialize the new disk:

   - If you have a backup of the previous disk's disk label information (Section 4.1.3):

     i. Apply the backup disk label to the new disk:

        # **disklabel -R** ***disk_access_name*** **auto** ***file***

     ii. Initialize the disk for LSM, using the disk access name:

        # **voldisk -f init** ***disk_access_name***

   - If no disk label file is available:

     i. Apply a default disk label to the new disk:

        # **disklabel -rwn** ***disk_access_name***

     ii. Initialize the disk for LSM:

        # **voldisksetup -i** ***disk***

g. Optionally (but recommended), create a backup copy of the new disk's disk label information:

   # **disklabel** ***disk_access_name*** **>** ***file***

h. Add the new disk to the applicable disk group, assigning a disk media name to the disk access name. You can reuse the disk media name of the failed disk as the disk media name for the new disk:

   # **voldg** [**-g** ***disk_group***] **-k adddisk**
   disk_media_name=disk_access_name

i. Verify that the volume's plex state has changed to RECOVER:

```
# volinfo -p
vol  tst           fsgen    Unstartable
plex tst-01        RECOVER
```

2. Reattach the log plex to the volume:

   ```
   # volplex att volume log_plex
   ```

6.4.6 Checking the Status of Volume Resynchronization

If the system fails and restarts, LSM automatically recovers all volumes that were running normally at the time of the failure.

- For volumes that use mirror plexes and have a DRL plex, this involves resynchronizing all the dirty regions.
- For volumes that use a RAID 5 plex and have a RAID 5 log plex, this involves replaying the log plex to complete any outstanding writes.

Configuring redundant volumes with log plexes is the recommended method to speed the recovery of volumes after a system failure. Under normal circumstances, the recovery happens so quickly that there is no noticeable effect (such as performance lag) once the system is running again. However, if the volume had no log, the resynchronization can take a long time (minutes to hours, or longer) depending on the size of the volume.

You can display the status of the volume resynchronization in progress to determine how long it will take. (You cannot check the status of *plex resynchronization*, which occurs when you replace a failed disk or add a new plex to a volume; the `volprint` command does not have access to that information. However, in these cases, the volume is usable while the resynchronization occurs.)

To determine the time remaining for a volume resynchronization in progress:

1. Check the read/write flags for the volume to see the current recovery offset value:

   ```
   # volprint -vl volume | grep flags
   ```

   Information similar to the following is displayed:

   ```
   flags:    open rwback (offset=121488) writeback
   ```

2. Check again after some time has passed (120 seconds is ample) to see how far the recovery has progressed:

   ```
   # sleep 120 ; volprint -vl volume | grep flags
   ```

   Information similar to the following is displayed:

   ```
   flags:    open rwback (offset=2579088) writeback
   ```

3. Calculate the rate of progress by dividing the difference between the offsets by the time that passed between the two checks. For example, in

120 seconds the resynchronization had completed 2457600 sectors. Each second, approximately 20480 sectors (10 MB) were resynchronized.

4. Multiply the resynchronization rate by the size of the volume, in sectors. This indicates the approximate amount of time a complete resynchronization will take. For example, at a rate of 20480 sectors per second, a volume that is 200 GB will take about five and a half minutes to resynchronize.

   The actual time required can vary, depending on other I/O loads on the system and whether the volume or the system experiences additional problems or failures.

6.4.6.1 Changing the Rate of Future Volume Resynchronizations

Although you cannot change the rate of (or stop) a volume resynchronization once it has begun, you can change the setting for the rate of future resynchronizations, if your volumes are large enough that the resynchronization has a noticeable impact on system performance during recovery.

Caution

Use this procedure only if you are a knowledgeable system administrator and you have evaluated the effect of volume resynchronization on system performance and determined it to be unacceptable. You must be familiar with editing system files and scripts.

To change the rate of volume resynchronization for future recoveries, use your preferred editor to modify the indicated line in the `/sbin/lsm-startup` script. The script contains information similar to the following, which has been edited for brevity and formatting:

```
#!/sbin/sh
   .
   .
   .
vold_opts=-k

volrecover_iosize=64k

s_flag=$1
   .
   .
   .

if [ "X`/sbin/voldctl mode 2> /dev/null`" = "Xmode: enabled" ]; then
     /sbin/volrecover -b -o iosize=$volrecover_iosize -s    [1]
       if [ $is_cluster -eq 1 -a $vold_locked -eq 1 ]
        then
                voldctl unlock
        fi
```

```
        if [ "$s_flag" != "-c" ]; then
                [ ! -f $STATEDIR/init_lsm ] && swapon -a > /dev/null 2>&1
        fi
        if [ "$s_flag" = "-c" ]; then
                Pid='/bin/ps -e | grep "volwatch" | awk '$6 != "grep"''
                if [ "X$Pid" = "X" ]; then
                        option='rcmgr -c get LSMSTART 2> /dev/null'
                        if [ "$option" = "mailplus" ]; then
                                /usr/sbin/volwatch -s &
                                egettxt "LSM volwatch Service started - \
                                         hot spare support" lsmshm.cat:5148
                        else
                                rcmgr -c set LSMSTART mailonly 2> /dev/null
                                /usr/sbin/volwatch -m &
                                egettxt "LSM volwatch Service started - \
                                         mail only" lsmshm.cat:5116
                        fi
                fi
                Pid='/bin/ps -e | grep "volnotify -e" | awk '$6 != "grep"''
                if [ "X$Pid" = "X" ]; then
                        volnotify_opts='rcmgr -c get LSM_EVM_OPTS 2> /dev/null'
                        if [ "$volnotify_opts" != "disable" ]; then
                                if [ ! -z "$volnotify_opts" ]; then
                                        /usr/sbin/volnotify -e $volnotify_opts > \
                                        /dev/null &
                                else
                                        /usr/sbin/volnotify -eicfd  >/dev/null &
                                fi
                        fi
                fi
        fi
else
        egettxt "LSM: Vold is not enabled for transactions" lsmshm.cat:981
        egettxt "  No volumes started\n" lsmshm.cat:982
        exit
fi
```

1 Change the indicated line to one of the following:

- To slow the rate of recovery, add `-o slow` as follows:

  `/sbin/volrecover -b -o iosize=$volrecover_iosize` **`-o slow`** `-s`

  This option inserts a delay of 250ms between each recovery operation. This can considerably reduce the performance impact on the system, depending on the size of the volume and the number of plexes.

- To disable resynchronization, add `-o delayrecover` as follows:

  `/sbin/volrecover -b -o iosize=$volrecover_iosize` **`-o delayrecover`** `-s`

  This option requires that you manually begin a resynchronization at your discretion, such as when the system is not under peak demand. Until then, the volume remains in read-writeback mode, which means that every time a region of the volume is read, the data is written to all plexes in the volume. When you eventually initiate the resynchronization, all regions marked dirty are resynchronized, perhaps unnecessarily.

This option incurs performance overhead by writing all reads back to all plexes, which might be less than the impact of permitting the resynchronization to complete during periods of high system demand.

You can change the `/sbin/lsm-startup` script back to its original state at any time.

6.4.7 Clearing Locks on LSM Volumes

When LSM makes changes to an object's configuration, LSM locks the object until the change is written. If a configuration change terminated abnormally, there might still be a lock on the object.

To determine if an object is locked, enter:

```
# volprint [-g disk_group] -vh
```

In the information displayed, the lock appears in the TUTIL0 column.

To clear the lock, enter:

```
# volmend [-g disk_group] clear tutil0 object ...
```

You might need to restart the volume. See Section 5.4.4.

6.5 Solving Disk Problems

The following sections describe troubleshooting procedures for failing and failed disks, including the boot disk.

6.5.1 Checking Disk Status

Disks can experience transient errors for a variety of reasons, such as when a power supply suffers a surge or a cable is accidentally unplugged. You can check the status of disks through the output of the `volprint` and `voldisk` commands.

To see the LSM status of a disk, enter:

```
# voldisk list
```

To check the usability of a disk, enter:

```
# voldisk check disk
```

Information similar to the following is displayed:

```
dsk5: Okay
```

The `voldisk` command validates the usability of the given disks by testing whether LSM can read and write the disk header information. A disk is considered usable if LSM can write and read back at least one of the disk headers that are stored on the disk. If a disk in a disk group is found to be unusable, it is detached from its disk group and all subdisks stored on the disk become invalid until you replace the physical disk or reassign the disk media records to a different physical disk.

Note

Because an LSM nopriv disk does not contain a disk header, a failed nopriv disk might continue to be considered okay and usable.

6.5.2 Recovering a Stale Subdisk

LSM usually recovers stale subdisks when the volume starts. However, it is possible that:

- The recovery process might get killed.
- The volume might be started with an option to prevent subdisk recovery.
- The disk on which the subdisk resides might have been replaced without any recovery operations being performed.

To recover a stale subdisk in a volume, enter:

```
# volume recover volume subdisk
```

To recover all stale subdisks in a volume, enter the same command without specifying a subdisk:

```
# volume recover volume
```

6.5.3 Recovering Volumes After a Temporary Disk Failure

If a disk had a temporary failure but is not damaged (for example, the disk was removed by accident, a power cable was disconnected, or some other recoverable problem occurred) and the system was not restarted, you can recover the volumes on that disk. (LSM automatically recovers volumes when the system is restarted.)

To recover from a temporary disk failure:

1. Make sure the disk is back on line and accessible; for example:
   - Check that the disk is firmly snapped into the bay.
   - Reconnect any loose cables.
   - Perform any other checks appropriate to your system.
2. Scan for all known disks to ensure the disk is available:

```
# voldctl enable
```

3. Recover the volumes on the disk:

   ```
   # volrecover -sb
   ```

6.5.4 Moving a Volume Off a Failing Disk

Often a disk has recoverable (soft) errors before it fails completely. If a disk is experiencing an unusual number of soft errors, move the volume off the disk and replace it.

Note

To replace a failed boot disk, see Section 6.5.6.1.

To move a volume off a failing disk:

1. Identify the size of the volume on the failing disk:

   ```
   # volprint [-g disk_group] -ht [volume]
   ```

2. Ensure there is an equal amount of free space in the disk group:

   ```
   # voldg [-g disk_group] free
   ```

   If there is not enough space, add a new disk. See Section 4.1.2.

3. Move the volume to a disk other than the failing disk, as specified by the `!` operand. Use the appropriate shell quoting convention to correctly interpret the `!`. You do not need to specify a target disk.

   ```
   # volassist [-g disk_group] move volume !disk
   ```

   See Section 6.5.5 for information on replacing a failed disk.

6.5.5 Replacing a Failed Disk

When an LSM disk fails completely and its state becomes DETACHED, you must:

1. Replace the disk with a new disk. For best results, replace a failed disk with the same or similar type of disk.
2. Recover the volumes that used the failed disk.
3. Optionally, if the failure caused data to be moved onto a spare disk, you can move the data onto the new disk.

Note

To replace a failed boot disk, see Section 6.5.6.1.

6.5.5.1 Replacing a Failed Disk

To replace a failed disk:

1. Identify the disk media name of the failed disk using one of the following commands:
   - To display all disk, disk group, and volume information, and the status of any volumes that are affected by the failed disk, enter:

     ```
     # volprint -Aht
     ```
   - To display only the disk information, enter:

     ```
     # volprint -Adt
     ```
2. Remove the failed disk and retain the disk media records:

   ```
   # voldg [-g disk_group] -k rmdisk disk_media_name
   ```
3. Remove the disk access records, using the disk access name:

   ```
   # voldisk rm disk_access_name
   ```
4. Physically remove the failed disk and replace it with the new disk.

   Please note that the device must be completely removed from LSM before running any non-LSM commands to remove and replace the failed disk, such as `hwmgr -redirect`.
5. Scan for the new disk:

   ```
   # hwmgr -scan scsi
   ```

   The `hwmgr` command returns the prompt before it completes the scan. You need to know that the system has discovered the new disk before continuing. See the `hwmgr`(8) reference page for more information on how to trap the end of a scan.
6. Label and initialize the new disk:
   - If you have a backup of the previous disk's disk label information (Section 4.1.3):
     a. Apply the backup disk label to the new disk:

        ```
        # disklabel -R disk_access_name auto file
        ```
     b. Initialize the disk for LSM, using the disk access name:

        ```
        # voldisk -f init disk_access_name
        ```
   - If no disk label file is available:
     a. Apply a default disk label to the new disk:

        ```
        # disklabel -rwn disk_access_name
        ```
     b. Initialize the disk for LSM:

```
# voldisksetup -i disk
```

7. Optionally (but recommended), create a backup copy of the new disk's disk label information:

   ```
   # disklabel disk_access_name > file
   ```

8. Add the new disk to the applicable disk group, assigning a disk media name to the disk access name.

   You can reuse the disk media name of the failed disk as the disk media name for the new disk. Use the `-k` option if you want to apply the existing LSM information for the failed disk to the new one:

   ```
   # voldg [-g disk_group] [-k] adddisk
   disk_media_name=disk_access_name
   ```

After you replace the disk, the steps you must do next, if any, depend on your setup:

- If hot-sparing occurred, the volume is running and requires no recovery.

  You can optionally move the data from the spare disk onto the disk you just replaced (Section 3.4.4.4), or you can configure the new disk as the hot-spare disk (Section 3.4.4.2).

- If hot-sparing did not occur, you must recover the volume or restore it from backup. See Section 6.5.5.2 for more information.

6.5.5.2 Recovering the Volumes

Use one of the following methods to recover the volume data:

- If the volume uses mirror plexes or a RAID 5 plex, start plex resynchronization. If the volume is large, you can run the resynchronization as a background task.

  ```
  # volrecover -sb volume
  ```

- If the volume is not redundant (not mirrored or RAID 5, or has no valid plexes from which to recover), restore the volume data from backup.

Optionally, verify the volume is started:

```
# volinfo
```

Information similar to the following is displayed:

```
home       fsgen    Started
finance    fsgen    Started
mkting     fsgen    Started
src        fsgen    Started
```

6.5.6 Recovering from a Boot Disk Failure

When the boot disk on a standalone system is encapsulated into an LSM volume with mirror plexes, failures occurring on the original boot disk are transparent to all users. However, during a failure, the system might:

- Write a message to the console indicating there was an error reading or writing to the plex on the boot disk
- Experience slow performance (depending on the problem encountered with the disk containing one of the plexes in the root or swap volumes)

To restart the system before you replace the original boot disk, you can boot from any disk that contains a valid `rootvol` volume.

If all copies of `rootvol` are corrupted and you cannot boot the system, you must reinstall the operating system.

Replacing a boot disk is a more complex process than replacing other disks because boot-critical data must be placed in specific areas on specific disks for the boot process to find it. How you replace a failed boot disk depends on:

- If you have mirrored the root disk and enabled hot-sparing support.
- If the errors are correctable and the same disk can be reused. This is known as readding a disk. If you reuse the boot disk, you should monitor it and replace it during your next maintenance cycle.
- If the disk has failed completely and must be replaced.

Section 6.5.6.1 gives instructions for replacing the boot disk, as well as other information related to boot disk recovery.

6.5.6.1 Replacing a Failed Boot Disk

The following procedure assumes that you originally encapsulated the boot disk on a standalone system and created mirror plexes for the boot disk volumes. The last step in this procedure creates a new (replacement) mirror on the new disk.

To replace a failed boot disk under LSM control with a new disk:

1. Restart the system from the disk that has not failed.
2. Display the status of all LSM disks and volumes to ensure you use the name of the failed disk and failed plex in the remaining steps:

   ```
   # voldisk list
   # volprint -ht
   ```

3. Dissociate the plexes on the failed disk from the root, swap, and user volumes, if `/usr` or `/var` were encapsulated on the boot disk.

   ```
   # volplex -o rm dis rootvol-02 swapvol-02 vol-dsk0g-02
   ```

The /usr and /var volumes have names derived from the partition letter of the boot disk (for example, vol–dsk0g).

4. Remove the failed LSM disks for the boot disk:

   a. Remove the disks from the rootdg disk group:

      ```
      # voldg rmdisk dskna dsknb dskng ...
      ```

   b. Remove the LSM disks configured on the boot disk from LSM control:

      ```
      # voldisk rm dskna dsknb dskng ...
      ```

5. Physically remove and replace the failed disk.

   Please note that the device must be completely removed from LSM before running any non-LSM commands to remove and replace the failed disk, such as `hwmgr -redirect`.

6. Scan for the new disk:

   ```
   # hwmgr -scan scsi
   ```

   The `hwmgr` command returns the prompt before it completes the scan. You need to know that the system has discovered the new disk before continuing. See the hwmgr(8) reference page for more information on how to trap the end of a scan.

7. Modify the device special files, reassigning the old disk name to the new disk. Make sure you list the new disk first.

   ```
   # dsfmgr -e new_name old_name
   ```

8. Label the new disk, setting all partitions to unused:

   ```
   # disklabel -rw new_disk
   ```

9. Mirror the existing root volumes onto the new disk:

   ```
   # volrootmir new_disk
   ```

The boot disk volumes are restored and ready for use.

6.6 Problems Importing a Disk Group

If you receive an error message when trying to import a disk group or the command fails, possible causes are:

- One or more of the disks contains the host ID of another system.

  To verify this, enter:

  ```
  # voldisk list disk_access_name
  ```

  If the host ID of the disk does not match that of the system where you are trying to import the disk group, enter:

```
# voldisk clearimport disk_access_name
```

You can now import the disk group.

- One or more of the disks might be inaccessible.

  Some disks might have failed. You can forcibly import the disk group and resolve the problem later; for example, replace the failed disk.

  To forcibly import a disk group, enter:

  ```
  # voldg -f import disk_group
  ```

  Once the disk group is imported, you can identify and solve the problem.

7

Error Messages

LSM is fault-tolerant and resolves most problems without system administrator intervention. If the volume daemon (`vold`) recognizes what actions are being taken, it can roll a transaction forward or back. If `vold` is unable to recognize and fix system problems, you need to solve the problem.

This appendix describes the majority of informational, failure, and error messages displayed by `vold` and the kernel driver. These sections include some errors that are infrequently encountered and difficult to troubleshoot. Causes are included to elaborate on the situation or problem that might have generated a particular message. Wherever possible, a recovery procedure (solution) is provided to locate and correct potential problems.

If you need to contact your customer support organization, these messages are numbered for ease of reference.

7.1 Volume Daemon Error Messages

The following list contains the error messages associated with the volume (`vold`) daemon.

1. `-r must be followed by 'reset'`

   Cause: This message is caused by a usage error.

   Solution: Correct the usage and try again.

2. `-x` *`argument`* `:  prefix too long`

   Cause: The stub-mode device path prefix name supplied exceeded the maximum of 32 characters.

   Solution: Select an alternate path for device files and retry the command.

3. `-x` *`string`* `:  invalid debug string`

   Cause: An unknown argument string was given to the `-x` option to `vold`.

   Solution: Select a valid string from the reference page for `vold` and try again.

4. ```
   Usage:  vold [-dkf] [-r reset] [-m mode] [-x level]
   For detailed help use:  vold help
   ```

   Cause: `vold` was invoked with an invalid set of arguments.
   ```

Solution: Correct the usage and try again or type `vold help` for more help. This is the full usage message from entering `vold help`:

```
Usage:  vold [-dkf] [-r reset] [-m mode] [-x level]
Recognized options:
 -d       set initial mode to disabled for transactions
 -k       kill the existing configuration daemon process
 -f       operate in foreground; default is background
 -r reset reset kernel state; requires 'reset' option argument
 -m mode  set vold's operating mode
          modes: disable, enable, bootload, bootstart
 -x level set debugging level to <debug>, 0 turns off debugging
 -R file  set filename for client request rendezvous
 -D file  set filename for client diag request rendezvous
```

5. `lsm:vold: Error: volume` *`volume`* `: Logging daemon killed by signal` *`signal_number`* `[ core dumped ]`

   Cause: Someone killed the logging daemon.

   Solution: If required, restart the daemon by entering `voliod logio`.

6. `lsm:vold: Error: /dev/volevent:` *`error_message`*
   `lsm:vold: Error: cannot open /dev/volconfig:` *`error_message`*
   `lsm:vold: Error: Cannot kill existing daemon, pid=` *`process_id`*

   Cause: An attempt to kill an existing `vold` process with a SIGKILL signal has failed. This might be due to the process being in an unkillable kernel state perhaps because of a hung I/O or a missing I/O interrupt. There might be disk driver error messages in the `/dev/osm` buffer.

   Solution: Try typing `cat /dev/osm` to see if any other messages have been sent to the console device. If possible, use `crash` to determine the state of the process. If the process is asleep waiting for an I/O completion, then any disk driver error messages that have occurred might point to the solution. Failing this, you should restart the system.

7. `lsm:vold: Error: /dev/voliod: VOL_LOGIOD_CHECK failed`

   Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

   Solution: Try stopping and restarting `vold`. If this fails then attempt to restart the system, possibly followed by reinstalling the LSM software. If this fails, contact Customer Support.

8. `lsm:vold: Error: /dev/voliod: VOL_LOGIOD_KILL failed`

   Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

Solution: Try stopping and restarting `vold`. If this fails then attempt to restart system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

9. `lsm:vold: Error: All transactions are disabled`

   Cause: This message might appear with the message `Disk group disabled by errors` if the disk group to be disabled is the root disk group. The continued use of the system could be dangerous because any configuration changes required (including error handling cases) could cause the loss of ability to perform I/O to a volume. Because this includes the root volume, this situation could, if uncorrected, cause the system to hang.

   Solution: This is a fatal error. All copies of the bootable root disk have failed. Recovery from this situation will require booting from floppy or from a disk unconnected with the LSM software. It might then be necessary to remove the LSM bootable disk configuration by using the `volunroot` command. See the LSM installation instructions for details. Then you can reinitialize the root disk group to reestablish the database and log areas.

10. `lsm:vold: Error: Cannot get all disk groups from the kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reinstalling the LSM software. If this fails, contact Customer Support.

11. `lsm:vold: Error: Cannot get all disks from the kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

12. `lsm:vold: Error: Cannot get kernel transaction state`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting the `vold` daemon. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

13. `lsm:vold:  Error:  Cannot get private storage from kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    *Solution:*Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

14. `lsm:vold:  Error:  Cannot get private storage size from kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

15. `lsm:vold:  Error:  Cannot get record` *`name`* `from the kernel:` *`error_message`*

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

16. `lsm:vold:  Error:  Cannot not make directory` *`directory_path`*

    Cause: When trying to create the specified directory, `vold` got a failure.

    Solution: Try creating the directory manually and then issue the command `voldctl enable`.

17. `lsm:vold:  Error:  Cannot recover operation in progress Failed to get group` *`disk_group`* `from the kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

18. `lsm:vold:  Error:  Cannot start` *`usage_type`* `volume, no valid complete plexes`

Cause: No usable plexes remain for either the root or swap volume. This error is fatal and will result in the system displaying the message `System startup failed` and then shutting down.

Solution: This is generally an unrecoverable error and will likely require that you reload the system from backups.

19. `lsm:vold: Error: Cannot start` *usage_type* `volume, no valid plexes`

    Cause No usable plexes remain for either the root or swap volume. This error is fatal and will result in the system displaying the message `System startup failed` and then shutting down.

    Solution: This is generally an unrecoverable error and will likely require a reload of the system from backups.

20. `lsm:vold: Error: Cannot start` *usage_type* `volume, volume state is invalid`

    Cause: The volume is not in a state that can be recovered from. This might be because of corruption of the databases or because of an invalid use of the `vold` interfaces without the use of the utilities.

    Solution: This is generally an unrecoverable error and will require reloading of the system from backups.

21. `lsm:vold: Error: Cannot store private storage into the kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

22. `lsm:vold: Error: Differing version of vold installed`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

23. `lsm:vold: Error: Disk` *disk*`, group` *disk_group*`, device` *device_name* `: not updated with new host ID Error:` *error_message*

    Cause: If the host ID for a system is changed using the `voldctl init` command then all disks in all imported disk groups will need to have the host ID changed to the new ID. If the host ID for a disk cannot be changed, then this message will be displayed. Other problems might also exist for this disk.

Solution: Move the contents of the disk elsewhere and reinitialize the disk.

24. `lsm:vold: Error: Disk group` *`disk_group`*`, Disk` *`disk`* `: Cannot auto-import group:` *`error_message`*

    Cause: The disk group *`disk_group`* could not be reimported after a system restart. The reason is included in the error message. Other error messages might appear that provide more information on what went wrong. Any volumes in the disk group will be unavailable until you fix the error condition and reimport the disk group.

    Solution: Clear the error condition, if possible, and then import the disk group manually with `voldg import`. After importing, restart all volumes with `voldg -g` *`groupname`* `-sb`.

25. `lsm:vold: Error: Disk group` *`disk_group`*`, Disk` *`disk`* `: Group name collides with record in rootdg`

    Cause: The disk group name *`disk_group`*, for the disk group being imported from the named disk, collides with a configuration record in the rootdg disk group. Disk groups must have names that do not match any records in the root disk group.

    Solution: If you want to import the disk group, rename the conflicting record in rootdg to some other name.

26. `lsm:vold: Error: Disk group` *`disk_group`* `: Cannot recover temp database error_message`

    Cause: The temp database stored in the root file system could not be opened or read. Other messages will detail the error. This might happen because of an I/O error or a problem in the file system.

    Solution: Restart the system and retry the operation.

27. `lsm:vold: Error: Disk group` *`disk_group`* `: Disabled by errors`

    Cause: This message can appear if the last configuration database or last kernel log area for a disk group became disabled. This could have been due to an I/O error or some other condition. Other messages preceding this one are likely to highlight the root cause.

    Solution: Back up any remaining active volumes. Reinitialize the disk group and add the disks back to the group to recover.

28. `lsm:vold: Error: Disk group` *`disk_group`* `: Errors in some configuration copies:`

    Cause: One or more on-disk database copies were found to contain errors. As a result, the disk group could not be imported. This is probably due to a disk I/O error, or to blocks of a configuration copy being overwritten within invalid contents. Check for messages from the disk driver. Errors pertaining to specific configuration copies are listed on successive lines. These lines can be in either of the following forms:

```
File filename : error_message : \
Block number : error_message
Disk diskname, copy copy_number : \
error_message : Block number : error_message
```

Lines beginning with `File` indicate an error in the special configuration copy file used for storing nonpersistent disk group information. Lines beginning with `Disk` indicate failure of a persistent configuration copy stored on a disk. The copy number indicates which of the disk's configuration copies contains the error.

Solution: If one or more disks for the disk group are currently inaccessible (such as due to a cabling error), make the disks accessible and try to import the disk group again with `voldg import`. Otherwise, the disk group is probably no longer usable and you will have to recreate it. All volume configuration information for the disk group is lost.

29. `lsm:vold: Error: Disk group` *`disk_group`* `: Reimport of disk group failed:` *`error_message`*

    Cause: The reload of a disk group into the kernel failed. This could be because the log size for the kernel might not be set or because of some other error in the import procedure. Other messages should indicate the true cause of the failure.

    Solution: Retry the operation unless some other error message leads to a suggested course of action. If this fails, restart the system.

30. `lsm:vold: Error: Disk group` *`disk_group`* `: update failed:` *`error_message`*

    Cause: This message occurs because a database update failed completely. No complete copy of the database could be written for the disk group. The disk group will be disabled and further access for configuration changes will be disallowed. If this error occurs for the root disk group, it will probably be necessary to reinstall the system.

    Solution: Back up any volumes still active in the disk group. Reinitialize the disk group and add the disks back to it.

31. `lsm:vold: Error: Exec of /sbin/voliod failed`

    Cause: An `exec` of `/sbin/voliod` failed.

    Solution: Check the existence and permissions of the `/sbin/voliod` command. Try executing the command manually to ensure that it can be run.

32. `lsm:vold: Error: Failed to store commit status list into kernel:`*`error_message`*

Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring of LSM. If this fails, contact Customer Support.

33. `lsm:vold: Error: Fork of logio daemon failed`

    Cause: The creation of a process that could then be used as a logging daemon failed.

    Solution: Check for messages explaining the reason that a `fork`(2) call failed. Retry the operation.

34. `lsm:vold: Error: GET_VOLINFO`
    `ioctl failed:` *`error_message`*
    `lsm:vold: Error: Version number of kernel does`
    `not match vold`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring of LSM. If this fails, contact Customer Support.

35. `lsm:vold: Error: Get of current rootdg failed`

    Cause: An attempt to retrieve the rootdg from the kernel failed. This might be because of a kernel `vold` inconsistency or could also be because of a version difference between `vold` and the kernel.

    Solution: Check that the correct version of `vold` and the kernel are installed. Other messages might suggest other problems in a prior attempt at loading a configuration and possible courses of action. If this fails, contact Customer Support.

36. `lsm:vold: Error: No convergence between root disk`
    `group and disk list Disks in one version of rootdg:`
    `disk` `type=` *`disk_type`* `info=` *`disk_info`* `Disks in alternate`
    `version of rootdg:` *`disk`* `type=` *`disk_type`* `info=`
    `disk_info`

    Cause: This message can appear when `vold` is not running in autoconfigure mode (see the `vold`(8) reference page) and when, after several retries, it cannot resolve the set of disks belonging to the root disk group. The algorithm for disks that are not autoconfigured is to scan disks listed in the `/etc/vol/volboot` file and then examine the disks to find a database copy for the rootdg disk group. The database copy is then read to find the list of disk access records for disks contained in the group. These disks are then examined to ensure that they contain the same database copy. As such, this algorithm expects to gain convergence on the set of disks and the database copies contained on them. If a loop is entered and convergence cannot be reached, then this message will appear and the root disk group importation will fail.

Solution: Reorganizing the physical locations of the devices attached to the system might break the deadlock. If this fails, contact Customer Support.

37. `lsm:vold: Error: Open of directory` *`directory_path`* `failed`

    Cause: When `vold` was trying to create node files for the volumes, it was unable to open the directory in which the nodes were to be created.

    Solution: Check for other errors that suggest why the directory might be missing or if the permissions might be incorrect. Fix the condition to allow `vold` to open or create the directory, then issue the command `voldctl enable`.

38. `lsm:vold: Error: Read of directory` *`directory_path`* `failed`

    Cause: The node directory could not be read when `vold` was trying to scan for volume nodes.

    Solution: Check for other messages that might suggest why the directory is inaccessible. Try reading the directory manually if the directory is corrupted, then try removing and recreating it and then restarting `vold`.

39. `lsm:vold: Error: Unexpected configuration tid for group` *`disk_group`* `found in kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring of LSM. If this fails, contact Customer Support.

40. `lsm:vold: Error: Unexpected error during` *`usage_type`* `volume reconfiguration:` *`error_message`*

    Cause: A record lock for the volume could not be acquired as part of the initial volume setup for either a root or swap volume. This is most likely to occur under low memory conditions.

    Solution: Other messages might suggest an alternate course of action. Otherwise, this is generally an unrecoverable error and will require that you either boot off an alternate root device or reload the system from backups.

41. `lsm:vold: Error: Unexpected error fetching disk for` *`usage_type`* `volume:` *`error_message`*

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring of LSM. If this fails, contact Customer Support.

42. `lsm:vold: Error: Unexpected values stored in the kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring of LSM. If this fails, contact Customer Support.

43. `lsm:vold: Error: VOL_RESET_KERNEL failed: a volume or plex device is open`

    Or:

    `lsm:vold: Error: VOL_RESET_KERNEL failed:` *`error_message`*

    Cause: An attempt at resetting the kernel state with a `vold -r reset` command failed because all the LSM objects in the kernel were not closed. If any volumes are in use, then the reset cannot be performed. This might also happen if a reset was requested on a system with root volumes. Root volumes are, by definition, never closed and so a reset cannot be performed.

    Solution: If you really want to reset the kernel, then check the state of the volumes and any mounted file systems to display information about who might have them open. Unmount all volumes and kill any processes accessing the volumes, then reset the kernel.

44. `lsm:vold: Error:` *`mode`* `: Unrecognized operating mode`

    Cause: An unknown mode string was entered following a `-m` option.

    Solution: Select a valid mode from the `vold`(8) reference page and try again.

45. `lsm:vold: Error: cannot open /dev/voliod:` *`error_message`*

    Cause: The open of the `/dev/voliod` file can fail only if the device node is missing or has an incorrect major or minor number.

Solution: Check the existence and values of the file and make sure that the LSM software was correctly installed.

46. `lsm:vold: Error: cannot open` *`argument`* `:` *`error_message`*

    Cause: The tracefile specified on the command line could not be opened in append mode. The error message supplied should explain the reason.

    Solution: Select an alternate tracefile name that can be created or appended to.

47. `lsm:vold: Error: cannot open` *`volconfig_device`* `: Device is already open`

    Or:

    `lsm:vold: Error: cannot open` *`volconfig_device`* `:` *`error_message`*

    Cause The exclusive open device (`/dev/volconfig`) is already open. Only one `vold` process can be active on the system at one time. Subsequent attempts at starting `vold` or opening the device will result in this message.

    Solution: Check for other running `vold` processes. The `voldctl mode` command will report if `vold` is currently active.

48. `lsm:vold: Error: enable failed:` *`error_message`*

    Cause: This message might occur during an initial startup of `vold`. If changing to enabled mode when this error occurs, failures could be due to problems with the creation of the portal or with connection to the kernel. If changing from an enabled state to a disabled state, then problems could occur with removing the disk groups from the kernel because of such things as volumes in use.

    Solution: Evaluate other error messages occurring with this one to determine the root cause of the problem. Make changes suggested by the other errors and then retry the command.

49. `lsm:vold: Error: failed to create daemon: fork failed:` *`error_message`*

    Cause: The call to `fork`(2) to generate a background `vold` process failed.

    Solution: Check for messages explaining the reason that a `fork`(2) call failed. Retry the operation.

50. `lsm:vold: Error: volume` *`volume`* `: Wait for logging daemon failed`

    Cause: The wait called to wait for the existence of the daemon process did not execute correctly. This can happen only if the `ioctl` does not correctly match the command required, perhaps because of a mismatch between the `voliod` command and the kernel versions or perhaps because of an incorrect minor number for the `/dev/voliod` device.

Solution: Check the existence and permissions of the `/dev/voliod` device.

51. `lsm:vold: FATAL Error: Disk group rootdg: Inconsistency -- Not loaded into kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring of LSM. If this fails, contact Customer Support.

52. `lsm:vold: FATAL Error: Group` *`disk_group`* `: Cannot update kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring of LSM. If this fails, contact Customer Support.

53. `lsm:vold: FATAL Error: Interprocess communication failure:` *`error_message`*

    Cause: The portal to client utilities has returned a failure. This is a fatal error because without a portal to clients, `vold` cannot do anything useful.

    Solution: Check for other errors suggesting the reason for portal failure. Restart `vold`. If problems persist, restart the system.

54. `lsm:vold: FATAL Error: Invalid status stored in kernel`

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring of LSM. If this fails, contact Customer Support.

55. `lsm:vold: Warning: Cannot create device` *`path`* `:` *`error_message`*

    Cause: The `mknod`(2) call made by `vold` to create a device node failed. The reason for the error should be displayed.

Solution: Fix the reason indicated for node creation failure and then issue the command `voldctl enable`.

56. `lsm:vold: Warning: Cannot exec /sbin/rm to remove` *`directory_path`* `:` *`error_message`*

    Cause: An `exec` of `/sbin/rm` failed.

    Solution: Ignore the error. It is not serious if the directory cannot be removed.

57. `lsm:vold: Warning: Cannot fork to remove directory` *`directory_path`* `:` *`error_message`*

    Cause: The call to `fork`(1) to generate a process that could then `exec` `rm`(2) failed.

    Solution: Ignore the error. It is not serious if the directory cannot be removed.

58. `lsm:vold: Warning: Disk` *`device_name`* `in kernel not a recognized type`

    Cause: The disk type of a disk in the kernel does not match any known disk type. This can occur only if `vold` and the kernel are in an inconsistent state.

    Solution: Try stopping and restarting `vold`. If this fails then reconfigure LSM. If this fails, contact Customer Support.

59. `lsm:vold: Warning: Disk` *`disk`* `names group` *`disk_group`*`, but group ID differs`

    Cause: As part of a disk group import, a disk was discovered that had a mismatched disk group name and disk group ID. This disk will not have been imported. This can happen only if two disk groups of the same name exist that have different disk group ID values. In that case, one group will be imported along with all its disks and the other group will not. This message will appear for disks in the unselected group.

    Solution: If it turns out that the disk should be imported into the group, add the disk to the group later. It will not happen automatically as part of the import. All configuration information for the disk will also be lost.

60. `lsm:vold: Warning: Disk group` *`disk_group`* `is disabled, disks not updated with new host ID`

    Cause: If the host ID for a system is changed using the `voldctl init` command then all disks in all imported disk groups will need to have the host ID changed to the new ID. If a disk group is found in the imported but disabled state, then the host ID will not be changed.

    Solution: Clear the host ID using the `voldisk clearimport` command for each disk, and then reimport the disk group.

61. `lsm:vold: Warning: Disk group` *disk_group* `: Disk group log may be too small Log size should be at least` *number* `blocks`

Cause: The log areas for the disk group have become too small for the size of configuration currently in the group. This should usually never happen without first displaying a message about the database area size. This message occurs only during disk group import; it occurs if the disk was inaccessible while new database objects were added to the configuration, and the disk was then made accessible and the system restarted.

Solution: If this situation does occur, then reinitialize the disks in the group with larger log areas. See the `voldisk`(8) reference page for more information. To reinitialize all the disks, detach them from the group with which they are associated, and then reinitialize and readd them to the disk group. Deport and reimport the group for the changes to the log areas for the group to take effect.

62. `lsm:vold: Warning: Disk group` *disk_group* `: Errors in some configuration copies:`

Cause: One or more on-disk database copies were found to contain errors. As a result, the disk group could not be imported. This is most likely to be due to a disk I/O error, or to blocks of a configuration copy being overwritten with invalid contents. Check for messages from the disk driver. Providing that other copies of the database can be successfully read, the system will continue and the disk group import or initial `vold` enable operation should succeed. If the database copy can subsequently be written to, then this message will not recur. Errors pertaining to specific configuration copies are listed on successive lines. These lines can be in either of the following forms:

```
File filename : error_message : Block number : error_message
Disk diskname, copy copy_number : error_message : \
Block number : error_message
```

Lines beginning with `File` indicate an error in the special configuration copy file used for storing nonpersistent disk group information. Lines beginning with `Disk` indicate failure of a persistent configuration copy stored on a disk. The copy number indicates which of the disk's configuration copies contains the error.

Solution: This message is likely to occur once due to an I/O failure and then not reoccur. If it does reoccur, then it might be necessary to remove the disk and reinitialize it to clear the condition. If all configuration copies for a disk group become unusable, then the disk group itself becomes unusable and must be recreated. If the rootdg disk group becomes unusable, you might need to reconfigure LSM. In this case, if the root file system is on a volume, then you might need to reinstall the operating system itself.

63. `lsm:vold: Warning: Error in volboot file:` *`error_message`* `Entry: disk` *`disk disk_type disk_info`*

    Cause: This message occurs when an entry in the `volboot` file does not contain the correct information to define a valid disk access record.

    Solution: Remove the entry using the `voldctl rmdisk` command and add it again using `voldctl adddisk`.

64. `lsm:vold: Warning: Failed to update voldinfo area in kernel:` *`error_message`*

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

65. `lsm:vold: Warning: Field too long in volboot file: Entry: disk` *`disk disk_type disk_info`*

    Cause: The `volboot` file is maintained by `vold` and `voldctl` and should never normally exhibit this problem. This problem might indicate some corruption of the `volboot` file or could also be the result of manual editing of the file.

    Solution: Try to remove the offending entry with the `voldctl rmdisk` command. If this fails, you might have to reinitialize `volboot` using the `voldctl init` command.

66. `lsm:vold: Warning: Get of record` *`record_name`* `from kernel failed:` *`error_message`*

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

67. `lsm:vold: Warning: Plex` *`plex`* `for` *`usage_type`* `volume is stale or unusable`

    Cause: This message is output to alert you to the failure of one or more plexes of either the root or swap volume. The system might be able to continue depending on the existence of other usable plexes for the volume.

Solution: Repair the failed plex by either reattaching the plex to the volume once the system is booted, or by evacuating and replacing the disk on which the failed plex resides if you think that the disk is going bad.

68. `lsm:vold: Warning: cannot remove group` *`group_id`* `from kernel:` *`error_message`*

    Cause: Some inconsistency between `vold` and the kernel has caused an `ioctl` to fail. This could be caused by the use of older versions of `vold` or the kernel, or it could be due to a bug in LSM.

    Solution: Try stopping and restarting `vold`. If this fails then restart the system, possibly followed by reconfiguring LSM. If this fails, contact Customer Support.

69. `lsm:vold: Warning: response to client` *`client number`* `failed:` *`error_message`*

    Cause: The portal to client utilities has returned a failure. This is a fatal error because without a portal to clients, `vold` cannot do anything useful. This could be caused by a STREAMS error or some other communications problem with the client.

    Solution: Check for other errors suggesting the reason for portal failure. Restart `vold`. If problems persist, restart the system.

7.2 Kernel Error Messages

The following are the kernel level error messages.

1. `NOTICE:` *`message`* `on volume device` *`hex_device_number`* `(` *`volume`* `) in diskgroup` *`disk_group`*

   Cause: This is caused by a driver above the LSM level calling the LSM `volprint` function. This usually happens when a driver detects some error condition in LSM and want to display the error.

   Solution: No action necessary, unless specified in a supplied string.

2. `NOTICE: io/vol.c(volerror): Correctable` *`type`* `error on volume` *`volume`*`, plex` *`plex`*`, block` *`block_number`*

   Cause: A correctable I/O error was detected and corrected. A correctable I/O error is one where a read error from an underlying device driver could be corrected by reading the data from an alternate mirror copy and then writing it back to the failed mirror.

   Solution: If the I/O could have been completed by reading from an alternate mirror but the writeback to the failed mirror still failed, the mirror will be detached. This failure will cause the exception handling code to recover the volume according to its error recovery policy. This usually results in either a mirror or the volume becoming detached.

You must reattach the mirror (`volplex att`), to bring back the failed mirror copy. If the volume was detached, then the data contained on it is unrecoverable and will have to be restored from backups.

3. `NOTICE: io/vol.c(volerror): Uncorrectable` *`type`* `error on volume` *`volume`*`, plex` *`plex`*`, block` *`block_number`*

   Cause: Following an I/O error from one mirror, an attempt to reread the data from an alternate mirror failed. This could be because no other mirrors exist or could be because the other mirrors also had I/O failures.

   Solution: This failure will cause the exception handling code to be entered, which will result in the volume's error recovery policy being followed. This can have effects ranging from detaching a mirror to disabling the volume. You must reattach the mirror (`volplex att`), to bring back the failed mirror copy. If the volume was detached, then the data contained on it is unrecoverable and you will have to restore it from backups.

4. `NOTICE: lsm: Can't open device` *`disk`*`, device busy or inaccessible.`

   Cause: The named disk cannot be accessed.

   Solution: Turn on the drive.

5. `WARNING: io/vol.c(volexcept): No volume error daemon - Cannot Log plex detach, detaching volume`

   Cause: No `voliod` process was running and able to log a detach record for a mirror that is being detached due to an I/O error. This is a fatal error that causes future access to the volume to be rejected, because any system failure coming after additional I/O would not be able to detect the failure of the mirror and mirror inconsistencies might then occur.

   Solution: Although it is too late to rescue this volume, you should start at least one `voliod` process as soon as possible (using `voliod set 2`). Stop and restart the failed volume, then restore the data from backups. Mirrors will have become inconsistent and so any attempt at using the data on the volume could prove disastrous.

6. `WARNING: volklog_dgfree: Can't clear group commit log record for group` *`disk_group`*

   Cause: This can occur if a log flush to disk could not be performed because no valid log copies remained. This is likely to compromise the ability of the LSM to recover from any further I/O errors.

   Solution: When you add disks to the system, ensure that new viable logging areas can be generated. Alternatively, remove failed disks and replace them with working devices.

7. `WARNING: volklog_dgfree:  Can't free kernel logging area for vol_reset_kernel of group` *`disk_group`*

   Cause: A free of the logs for a disk group failed because either no valid log areas remained for flushing or some log records remained in the log before the clear operation was requested.

   Solution: There is no action to take here; this is an LSM internal error. Contact Customer Support.

A

The Storage Administrator (lsmsa)

This chapter describes how to manage AdvFS domains, filesets, and volumes and LSM objects, including disks, disk groups, volumes, plexes, and subdisks using the Storage Administrator GUI, also called `lsmsa`. The tasks described in this chapter can also be accomplished by using:

- The command line.
- The `voldiskadm` menu interface. The `voldiskadm` interface does not administer AdvFS. See Appendix C for more information on the `voldiskadm` menu interface.
- The Visual Administrator, also known as `dxlsm`. This interface does not administer LSM in a cluster. See Appendix D for more information on the Visual Administrator.

See Appendix B for more information on how to track Storage Administrator activities and how to customize the Storage Administrator GUI.

A.1 Overview

The Storage Administrator is a Java-based graphical user interface (GUI) for LSM. The Storage Administrator displays a hierarchical view of LSM objects, AdvFS domains, and their relationships. You use the Storage Administrator to view and manage LSM objects and AdvFS domains on a local or remote (client) system.

The Storage Administrator consists of a server (daemon) and a client. The Storage Administrator server runs on a system that you want to administer. If LSM is not initiated, then the interface starts up in AdvFS-only mode. The Storage Administrator client runs on any machine that supports the Java run-time environment.

The Storage Administrator provides dialog boxes in which you enter information to create or manage LSM objects or AdvFS domains. Completing a dialog box might be the equivalent of entering several commands.

Note the following considerations when using the Storage Administrator:

- If you are working in a TruCluster environment, some restrictions apply. Dialog box options for invalid tasks are grayed out.

- Mirrors are allowed in clusters except for the individual member boot partitions and the quorum disk.
- Software-based RAID 5 technology is not supported in clusters. If you want RAID 5 functionality, you must use hardware RAID devices.
- AdvFS file systems are supported in all modes.
- UFS file systems are supported in read-only mode.

A.1.1 Installing and Starting the Storage Administrator GUI

To install the Storage Administrator, choose the LSM GUI option during the LSM installation. Install the Storage Administrator on all systems on which you want to use the Storage Administrator to remotely manage the LSM software. See Chapter 3 for information on initializing the LSM software.

To use the Storage Administrator, you must log in as `root` unless your user name is in the `/etc/group` file for the system, in the group defined for Storage Administrator administration (`lsmsa_admin` by default).

To start the Storage Administrator, enter:

```
# /usr/bin/lsmsa
```

The Session Initiation dialog box is displayed. Enter the following information, then click on Ok:

- The name of the system to be administered in the Host Server field
- `root` or your user name
- The password associated with the account

When the Storage Administrator is started, the system attempts to connect it with the server process on the indicated host.

- If the Storage Administrator cannot connect to the server process, the system then attempts to connect it to the indicated host at the `initlsmsad` port, defined in the `/etc/services` file, and `vrts.remote.server.initLsmsadPort` port defined in the `/usr/lib/java/applications/lsmsa/properties` file. The port number defined in these two files must be the same.

  When the Storage Administrator connects to the `initlsmsad` port, the `inetd` server executes the `/usr/lib/java/applications/lsmsa/initlsmsad` program, which creates a subprocess where the `/usr/sbin/lsmsad` script will run. The `lsmsad` script starts the LSMSA server processes `VMServerImpl`, `VRTSRegistry` and `cmdserver`. Once the LSMSA server processes are started, the Storage Administrator then connects and operates normally.

- If the Storage Administrator cannot connect to a port, the error message "`Cannot connect to the server`" is displayed. When the Storage Administrator exits and disconnects from the server, the server continues to exist in an idle state until another GUI connects or an LSM configuration event occurs (such as creation, deletion, or modification of an LSM object). If the server receives notification of an LSM configuration event and if no GUIs are connected, the server exits. When the server exits, all the Storage Administrator server processes exit as well.

If the Storage Administrator cannot connect to the server, try the following:

- Check the `/var/lsmsa/logs/server.log` file for startup and error messages.
- Run the `/usr/lib/java/applications/lsmsa/initlsmsad` program to view error messages. You must be `root` user.
- On a very slow network, you might need to adjust the value assigned to the `CONNECTION_TIMEOUT` variable in the `/usr/sbin/lsmsad` script. This value is the amount of time after startup that the `lsmsa` server will wait for a connection from the client. The default value is 30 seconds. If the server process receives an LSM configuration event and if no GUIs have been or are currently connected, and if the `CONNECTION_TIMEOUT` seconds have elapsed, the server exits.

A.1.2 Storage Administrator Main Window Components

The Storage Administrator main window consists of two panes. The left pane displays a hierarchical tree of objects. The right pane displays an object table that lists the properties of the object you select on the left. The window also has a Menu Bar and a Command Launcher that you can hide or display to initiate LSM options. Figure A–1 shows the components of the main window.

Figure A–1: Storage Administrator Main Window

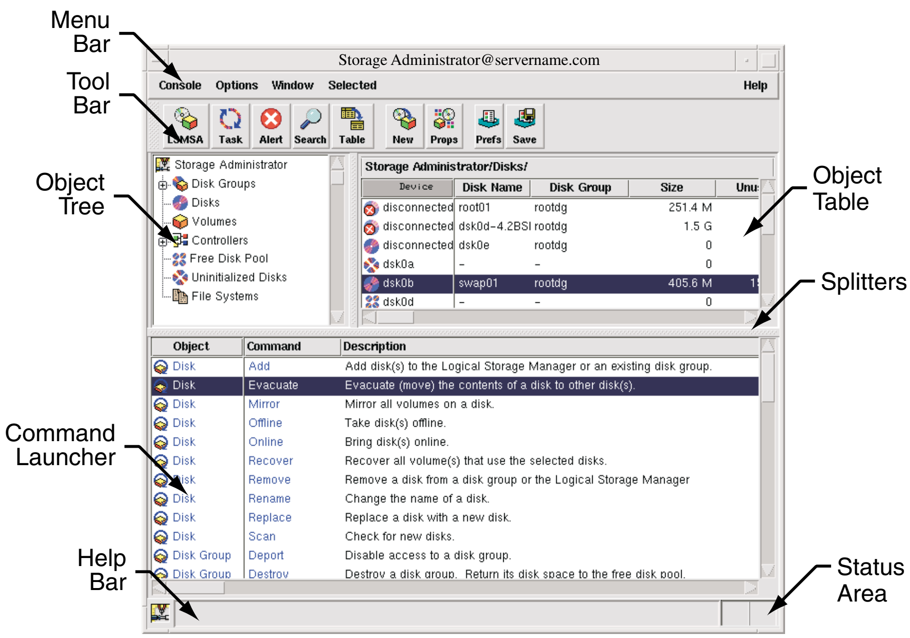

The Object Tree displays the hierarchical relationship between LSM objects and between objects in AdvFS domains. Each object represents a group of components of the same type. When you select the icon to the left of an object, components of that type appear in the right pane. You can expand objects (by clicking on the plus sign) to display their hierarchy.

The Object Table displays component objects that belong to the currently selected object in the Object Tree. The Object Table is dynamic and constantly updates its contents to reflect changes to the system.

The Command Launcher displays a list of tasks that you can perform on objects. When you click on a task in the Command Launcher list, the task starts and the dialog box for the task appears.

The Menu Bar contains the following menus:

- Console menu—Provides access to the New menu, which creates volumes, disk groups, or file systems. It also closes the Main window, provides access to an object properties dialog box, or exits the Storage Administrator.
- Options menu—Provides access to the Preferences dialog box and saves or loads user preferences for Storage Administrator components. The Options menu also removes alert icons from the Status Area.

- Window menu—Opens additional Storage Administrator Main windows, the Task Request Monitor, the Alert Monitor window, the Search window, a copy of the Object Table, or the Command Launcher.
- Selected menu—By default, the Selected menu is grayed out and changes its options based on the type of object that you select. Until you have selected an object, you see the grayed Selected entry on the Menu Bar. When you select an object, for example Volumes, the Selected entry changes to Volumes.
- Help menu—Provides access to online Storage Administrator help.

The toolbar consists of the following buttons that provide access to the following windows:

- LSMSA button—Launches an additional Storage Administrator Main window.
- Task button—Launches the Task Request Monitor window.
- Alert button—Launches the Alert Monitor window.
- Search button—Launches the Object Search window.
- Table button—Launches a window that contains a copy of the main Object Table.
- New button—Launches the New Volume dialog box that is used to create a volume.
- Props button—Launches the Object Properties dialog box for a selected object.
- Prefs button—Launches the Preferences dialog box.
- Save button—Saves the current preference settings for use in future Storage Administrator sessions.

A.2 Performing Tasks with the Storage Administrator

You perform most tasks by selecting objects or tasks, then providing information in resulting dialog boxes. You perform tasks using:

- The Menu Bar
- A pop-up menu
- The Command Launcher

A.2.1 Selecting Objects

To select a single object, click on it. To deselect the object, click on it again.

To select or deselect multiple objects, hold down the Control key while selecting the objects. The objects that you select do not have to be adjacent.

To select a range of adjacent objects, select the first object and then hold down the Shift key while selecting the last object in the range. You can also select multiple adjacent objects by dragging the mouse over the desired objects while pressing the Shift key.

A.2.2 Using the Console Menu and Selected Menu

You can launch tasks from the Console and Selected menus in the Menu Bar.

Choose New from the Console menu to create:

- LSM volumes and disk groups
- File systems
- AdvFS domains, volumes, and filesets

The context-sensitive Selected menu changes to reflect which object you selected in the Object Table.

For example, to change a volume name, select Volumes in the Object Tree and the volume you want to rename in the Object Table. From the Volumes menu choose Rename, enter information in the Rename Volume dialog box, and click on OK.

A.2.3 Using the Right Mouse Button to Display a Pop-Up Menu

Click on the right mouse button to access a context-sensitive pop-up menu to display common task information that you can apply to the selected object. Additional tasks are available through the Menu Bar or the Command Launcher.

For example, to create a new volume in a disk group, select Disk Groups in the Object Tree, right-click on the disk group in the Object Table, choose New Volume from the pop-up menu, enter information in the New Volume dialog box, and click on OK.

A.2.4 Using the Command Launcher Window

The Command Launcher window contains a list of objects and associated tasks. To display the Command Launcher window, choose Command Launcher from the Window menu. To hide the Command Launcher window, choose Command Launcher again from the Window menu.

To perform a task on a specific type of object, select the appropriate object-command combination from the Command Launcher list. For example, to create a volume, choose Volume–New from the Command Launcher, enter information in the New Volume dialog box, and click on OK.

Caution

The Command Launcher does not restrict context. If you choose inappropriate commands and ignore the warning dialogs, you can perform operations that might result in permanent loss of data.

A.2.5 Exiting the Storage Administrator

To exit the Storage Administrator, choose Exit from the Console menu.

A.3 Managing LSM Disks

The following sections describe LSM disk management tasks that you can complete using the Storage Administrator.

A.3.1 Adding an Entire Disk

When you add a disk for use with the LSM software, the disk is either initialized or encapsulated. If the disk is not set up, initialize it. If you want to use a disk with partitions that are in use, encapsulate it. Encapsulation preserves any existing data on the disk in the form of volumes. Initialization destroys any existing data on the disk. Initialized disks are placed in the free disk pool and are available to add to a disk group.

To add a disk for use with the LSM software:

1. In the Object Tree, select Disks and in the Object Table, select a disk to place under LSM control.

   If the disk is new and does not show up in the Object Table when you highlight Disks, click on Storage Administrator (at the top of the Object Tree) and choose Scan Disks from the System menu.

   In the Scan Disks dialog box, click on Ok to begin the search. From the Object Tree displayed, select a disk to add.

2. From the Disks menu, choose Add.

   The Add Disk(s) dialog box (Figure A–2) is displayed.

Figure A–2: Add Disk(s) Dialog Box

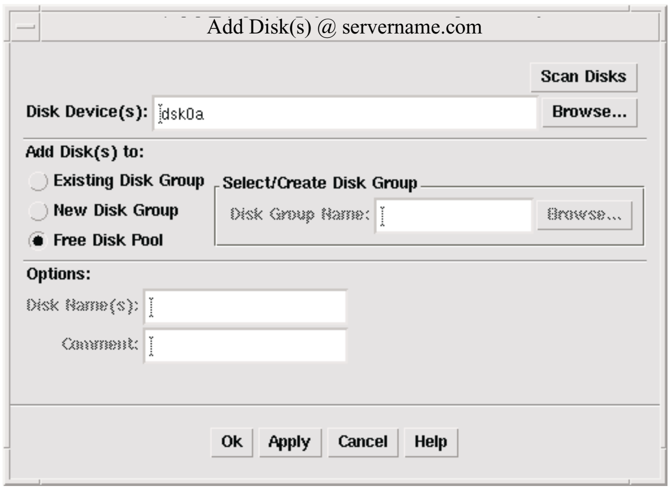

3. In the Add Disk(s) dialog box:

    a. If the correct disk device name is not displayed in the Disk Device field, enter the disk device name or click on Browse to select the disk. You can enter more than one name separated by spaces.

    b. Specify where to add disks:

        - To add disks to an existing disk group, select Existing Disk Group. Enter the disk group name in the Disk Group Name field or click on Browse to select a disk group.
        - To add disks to a new disk group, select New Disk Group. Enter the name of the new disk group in the Disk Group Name field. The new disk group is created.
        - To place disks in the free disk pool, select Free Disk Pool. Disks in the free disk pool are under LSM control (initialized) but do not belong to a disk group and cannot be used to create volumes.

    c. In the Options section:

        i. Specify the LSM disk names for the disks by typing a disk name in the Disk Name(s) field. This name must be unique within the disk group. If no LSM disk name is specified, the Storage Administrator assigns a default name to the disk.

        ii. Enter a comment if desired.

    d. Click on Ok.

A.3.2 Adding a Hot-Spare Disk

You can add one or more unused disks to a disk group and designate them as hot-spare disks. If an I/O failure occurs, the hot-spare feature automatically relocates any redundant (mirrored or RAID 5) subdisks to the spare disk and restores the affected LSM objects and data. You are notified of the failure and relocation details by electronic mail. See Section 3.4.4 for more information on the hot-spare feature.

If you designate a hot-spare disk, provide at least one per disk group. In the event of disk failure, the hot-spare disk automatically replaces the failed disk. Volumes can use hot-spare disks only from within the same disk group.

To add a disk as a hot-spare disk:

1. In the Object Tree, select Disk Groups and in the Object Table, select the LSM disk to be designated as a hot-spare disk.
2. From the Disk Groups menu, choose Properties.
3. In the Disk Properties window:
   a. Select the General tab.
   b. Select Spare.
   c. Click on Ok.

Note

The Properties dialog box is associated with the disk you have highlighted when you choose Properties. If you select a different disk from the Object Table, you must open a new Properties dialog box.

A.3.3 Evacuating a Disk

You can evacuate (or move) the contents of a disk to other disks in the same disk group if there is sufficient free space. If no target disk is specified, LSM uses available disks with sufficient free space. Evacuating a disk is useful in the event of disk failure.

If the disk being evacuated contains part of a mirrored, striped, or RAID 5 volume, do not move the contents of the disk to another disk containing a copy of the mirrored volume or part of the striped or RAID 5 volume.

To evacuate a disk:

1. In the Object Tree, select Disks and in the Object Table, select the disk that contains the objects and data to be moved.

2. From the Disks menu, choose Evacuate.

   The Evacuate Disk dialog box (Figure A–3) is displayed.

Figure A–3: Evacuate Disk Dialog Box

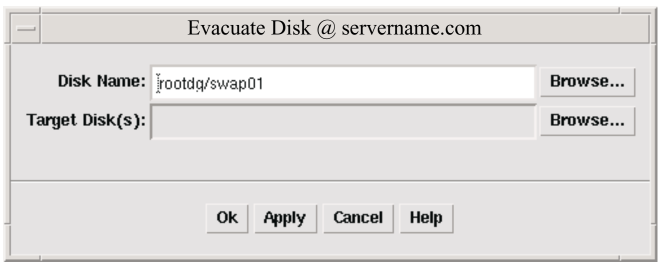

3. In the Evacuate Disk dialog box:

   a. If the correct disk name is not displayed in the Disk Name field, enter the disk name or click on Browse then click on the Object Tree to select the disk.

   b. Enter the name of the target disk to which you want to move the contents of the evacuated disk or click on Browse then click on the Object Tree to select one or more target disks.

      If you choose Browse, the total evacuated space (in kilobytes) is displayed so you can choose your target disk accordingly.

   c. Click on Ok.

A.3.4 Mirroring a Disk

You can mirror all LSM objects on a disk to provide high availability for the volumes on that disk. Mirroring also improves read performance, because multiple reads to the same volume can be done simultaneously using the multiple copies of data.

If possible, mirror the objects onto a disk on a different bus to reduce the risk of a single point of failure for the volumes.

To mirror all volumes on a disk:

1. In the Object Tree, select Disks and in the Object Table, select the disk that contains the volumes to be mirrored onto another disk.

2. From the Disks menu, choose Mirror.

   The Mirror Disk dialog box (Figure A–4) is displayed.

Figure A–4: Mirror Disk Dialog Box

3. In the Mirror Disk dialog box:

   a. If the correct disk name is not displayed in the Disk Name field, enter the disk name or click on Browse to select the disk.

      If you choose Browse, the mirrored total space is displayed so you can choose your target disk accordingly.

   b. To specify the disks to contain the new mirrors, enter the target disk name or click on Browse and complete the Target Disk dialog box.

   c. Click on Ok.

A.3.5 Taking a Disk Off Line

You can take a disk off line to prevent LSM from accessing it. You must remove a disk from its disk group before you take it off line. An offline disk remains unavailable until you restore access to the disk by placing it on line.

You take a disk off line to protect it from unintentional use; for example, if attempts to access it might have a negative effect on the system. You cannot take a disk that is in use off line.

To take a disk off line:

1. In the Object Tree, select Disks and in the Object Table, select the disk to be taken off line.
2. From the Disks menu, choose Offline.

   The Offline Disk dialog box is displayed.
3. In the Offline Disk dialog box:

   a. If the correct disk name is not displayed in the Disk Name field, enter the disk name or click on Browse to select the disk.

   b. Click on Ok.

A.3.6 Placing a Disk On Line

Placing a disk on line restores access to a disk that is off line. The disk is placed in the free disk pool and is accessible to LSM again. After bringing a disk back on line, the disk must be added to a disk group before it can be used for volumes.

Only disks that are off line can be placed on line.

To place a disk on line:

1. In the Object Tree, select Disks and in the Object Table, select the disk to be brought on line.
2. From the Disks menu, choose Online.

   The Online Disk dialog box is displayed.
3. In the Online Disk dialog box:

   a. If the correct disk name is not displayed in the Disk Name field, enter the disk name or click on Browse to select the disk.
   b. Click on Ok.

A.3.7 Recovering Volumes on a Disk

A recovery operation depends on the types of volumes on the disk and includes starting disabled volumes, resynchronizing mirrors in mirrored volumes, and resynchronizing parity in RAID 5 volumes.

Alert icons and the Alert Monitor window might indicate when a volume recovery is needed.

If recovery of a volume is not possible, restore the volume from backup.

To recover all volumes on a disk:

1. In the Object Tree, select Disks and in the Object Table, select the disk that contains the volumes to be recovered.
2. From the Disks menu, choose Recover.

   The Recover Disk dialog box is displayed.
3. In the Recover Disks dialog box:

   a. If the correct disk name is not displayed in the Disk Name field, enter the disk name or click on Browse to select the disk.
   b. Click on Ok.

A.3.8 Removing an LSM Disk from a Disk Group

An LSM disk no longer in use can be removed from a disk group.

Caution

Do not remove LSM disks that are currently in use (for example, disks that contain subdisks for a volume); doing so can result in loss of data or of data redundancy.

After an LSM disk is removed from a disk group, it is still initialized for use with the LSM software. Therefore, after removing the disk from a disk group, it can be either immediately added to another disk group, removed from LSM, or left for later use.

To remove an LSM disk from a disk group:

1. In the Object Tree, select Disks and in the Object Table, select the disk to be removed.
2. From the Disks menu, choose Remove.

   The Remove Disk dialog box (Figure A–5) is displayed.

   Figure A–5: Remove Disk Dialog Box

   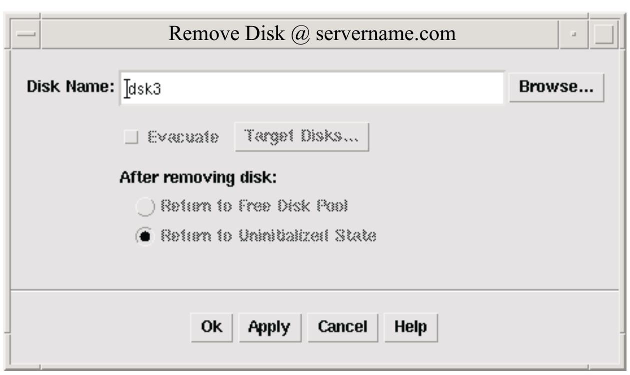

3. In the Remove Disk dialog box:

   a. If the correct disk name is not displayed in the Disk Name field, enter the disk name or click on Browse to select the disk.

   b. To move the contents of the disk to another disk before the disk is removed, select Evacuate. Click on Target Disks to specify one or more disks to which you want the contents moved.

   c. Specify how to handle the disk after removal:

- To remove the disk from its disk group and place it in the free disk pool, select Return to Free Disk Pool. The disk remains under LSM control.
- To remove the disk from LSM control, select Return to Uninitialized State.

d. Click on Ok.

A.3.9 Renaming a Disk

Because disk access names are defined by the operating system and media names are defined by you, you can rename only disk media names for disks in a disk group.

To rename the disk media name for an LSM disk:

1. In the Object Tree, select Disks and in the Object Table, select the disk to be renamed.
2. From the Disks menu, choose Rename.

   The Rename Disk dialog box is displayed.
3. In the Rename Disk dialog box:

   a. If the correct disk name is not displayed in the Disk Name field, enter the disk name or click on Browse to select the disk.

   b. Enter the new LSM disk name.

   c. Click on Ok.

A.3.10 Replacing a Disk

You might need to replace a disk if the disk fails and needs to be removed and repaired. You can replace an existing disk with a new physical disk, move volumes to the new disk, and attempt to recover any redundant (mirrored or RAID 5) volumes on the disk. You cannot recover nonredundant volumes. You should restore nonredundant volumes from backup. If the disk being replaced is a boot disk, you can set up the new disk as a boot disk.

If you replace a good disk, you need to remove the disk from its disk group before you replace the disk. If you replace a disk that has failed and is disconnected, you do not need to remove the disk from the disk group.

To replace a disk:

1. In the Object Tree, select Disks and in the Object Table, select the disk to be replaced.
2. From the Disks menu, choose Replace.

The Replace Disk dialog box (Figure A–6) is displayed.

Figure A–6: Replace Disk Dialog Box

3. In the Replace Disk dialog box:

   a. If the correct disk name is not displayed in the Disk Name field, enter the LSM disk name for the disk to be replaced or click on Browse to select the disk.

   b. Enter the physical disk name for the new (replacement) disk or click on Browse to select a disk.

   c. Click on Ok.

A.3.11 Scanning for New Disks

You can search your configuration for disks that are not under LSM control. Disks that are found are added to the free disk pool.

To scan for a new disk:

1. Select a disk from the Object Table.
2. From the Disks menu, choose Add.
3. Click on Scan Disk.
4. To view disks that are found, click on Free Disk Pool in the Object Tree.

A.4 Managing Disk Groups

The following sections describe disk group management tasks that you can complete using the Storage Administrator.

A.4.1 Creating a Disk Group

You must place disks into a disk group before you can use them to create volumes. The default disk group (rootdg) is created during LSM installation and always exists on a system running LSM. You can create additional disk groups to organize your disks into logical sets.

Each new disk group must contain at least one disk and must have a unique name. You can use only disks that are on line and do not already belong to a disk group.

To create a disk group:

1. From the Console menu, choose New then Disk Group.

   The New Disk Group dialog box (Figure A–7) is displayed.

Figure A–7: New Disk Group Dialog Box

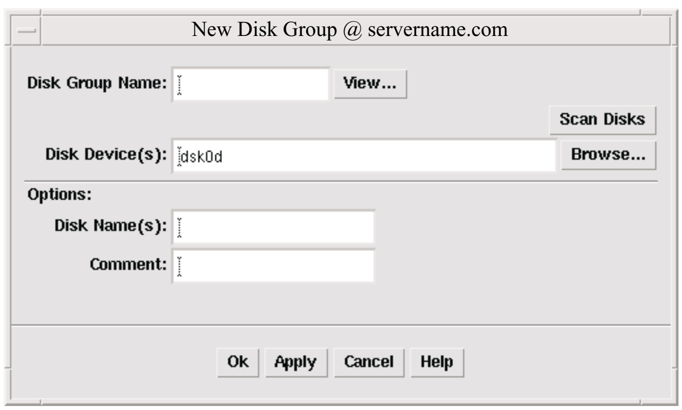

2. In the New Disk Group dialog box:

   a. Enter the name of the disk group to be created. Click on View to view the names of existing disk groups.

   b. To set up any new disks on the system, click Scan Disks. This runs the disk setup commands appropriate for the operating system.

   c. Enter the disk devices to be placed in the new disk group or click on Browse to select the devices.

   d. There are two options:

      - To specify the LSM disk name for the disk, enter a disk name in the Disk Name(s) field. If no LSM disk name is specified, the Storage Administrator assigns a default name to the disk.

      - Enter a comment if desired.

   e. Click on Ok.

A.4.2 Adding a Disk to a Disk Group

To add a disk to a disk group, follow the instructions for adding a disk (Section A.3.1). In the Add Disk dialog box, specify an existing disk group.

The LSM disk name must be unique within the disk group. If multiple disks are specified in the Disk Device(s) field and only one disk name is specified in the Disk Name(s) field, LSM appends numbers to the disk name so that each disk name is unique within its disk group.

You must place disks that belong to a disk group in the free disk pool before you can add them to another disk group. You must add disks in the free disk pool to a disk group before you can use them to create volumes.

Disks must be on line before you can add them to a disk group or the free disk pool. You cannot add disks to deported disk groups.

You must place the root disk in the root disk group (rootdg). If the root disk is placed in any other disk group, you cannot use the root disk to boot the system.

A.4.3 Adding a Disk Partition to a Disk Group

To add a disk partition to an LSM disk group:

1. In the Object Tree, select Free Partitions and in the Object Table, select the partition to be added.
2. From the Free Partitions menu, choose Add.

   The Add Disk Partition dialog box (Figure A–8) is displayed.

Figure A–8: Add Disk Partition Dialog Box

3. Enter the name of the Disk Group to which you will add a partition or click on Browse to select a name.
4. Click on Ok.

A.4.4 Deporting a Disk Group

After a disk group is created, the LSM software automatically imports it for use whenever the system is booted.

To disable access to a disk group, you deport the disk group. You must stop all the volumes within the disk group before deporting it.

To deport a disk group:

1. In the Object Tree, select Disk Groups and in the Object Table, select the disk group to deport.
2. From the Disk Groups menu, choose Deport.

   The Deport Disk Group dialog box (Figure A–9) is displayed.

Figure A–9: Deport Disk Group Dialog Box

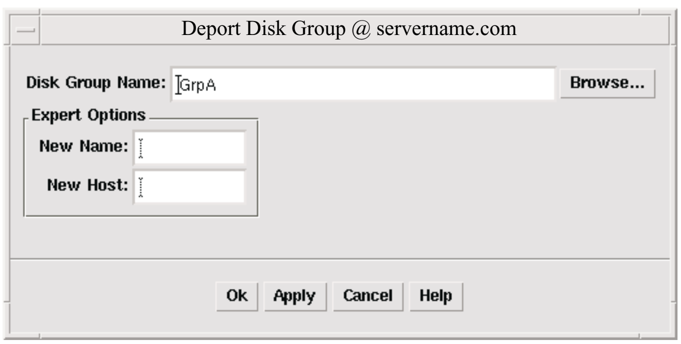

3. In the Deport Disk Group dialog box:

   a. If the correct disk group name is not displayed in the Disk Group Name field, enter the disk group name or click on Browse to select the disk group.

   b. Use the following Expert Options with caution:

      - To change the name of the disk group at deport, enter a new disk group name in the New Name field.
      - To set up a host machine to import the deported disk group at restart, enter the host ID in the New Host field.

   c. Click on Ok.

A.4.5 Importing a Deported Disk Group

You can import a disk group to make a deported (inaccessible) disk group and its volumes accessible again. To import a deported disk group, you must know the disk group's former name. This disk group name must have remained unused. In addition, at least one disk formerly assigned to the deported disk group must remain unused. If all disks associated with a deported disk group were reused because the disk group was deported, that disk group cannot be imported.

The import might fail for a number of reasons; for example, if the host cannot find one or more disks in the disk group. If the import fails because a disk has failed, you can import the disk group by selecting the Force

Import Expert option. If the import fails for another reason, a forced import can cause serious problems.

When you import a disk group, the system stamps its host ID on all disks in the disk group. A disk group import fails if one of the disks is stamped with a host ID that does not match the others. This ensures that dual-ported disks cannot be managed (and possibly corrupted) by two systems at the same time. If you are sure that the disk group is not in use by another host, you can clear the host IDs and import the disk group by selecting the Clear Host ID Expert option.

To import a deported disk group:

1. In the Object Tree select Disk Groups.
2. From the All Disk Groups menu, choose Import Disk Group.

   The Import Disk Group dialog box (Figure A–10) is displayed.

Figure A–10: Import Disk Group Dialog Box

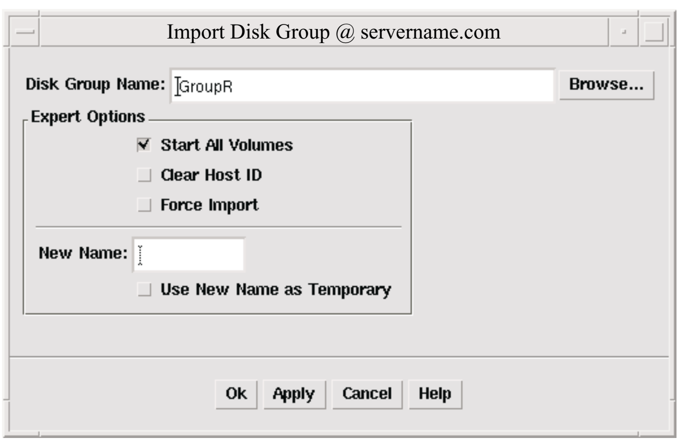

3. In the Import Disk Group dialog box:

   a. If the correct disk group name is not displayed in the Disk Group Name field, enter the disk group name or click on Browse to select the disk group.

   b. Use the following Expert Options with caution:

      - To start all volumes in the disk group at import, select Start All Volumes.
      - To clear the existing host ID stamp on all disks in the disk group at import, select Clear Host ID. Do not use this option if another host is using any disk in the disk group.

- To force the disk group import when the host cannot access all disks in the disk group, select Force Import.

c. Enter the name of the disk group to be imported or click on Browse to select the disk group.

d. To change the name of the disk group at import, enter a new disk group name in the New Name field. To indicate that the name change is temporary, select Use New Name as Temporary. If you indicate a temporary name change, the original name is restored when the system is restarted.

e. Click on Ok.

A.4.6 Recovering Volumes in a Disk Group

You can recover volumes in a given disk group. The recovery operations depend on the types of volumes in the disk group and include starting disabled volumes, resynchronizing mirrors in mirrored volumes, and resynchronizing parity in RAID 5 volumes.

Alert icons and the Alert Monitor window might indicate when volume recovery is needed.

In some cases, recovery might not be possible. If the volume recovery fails, you can attempt to restore the volume from backup.

To recover all volumes in a disk group:

1. In the Object Tree, select Disk Groups and in the Object Table, select the disk group containing the volumes to be recovered.
2. From the Disk Groups menu, choose Recover.

   In the Recover Disk Groups dialog box, click Yes to recover the volumes in the disk group.

A.4.7 Renaming a Disk Group

You can rename a disk group. The new disk group name must be unique. If volumes in the disk group are in use (mounted), the disk group cannot be renamed.

Renaming a disk group updates the `/etc/fstab` file.

To rename a disk group:

1. In the Object Tree, select Disk Groups and in the Object Table, select the disk group to be renamed.
2. From the Disk Groups menu, choose Rename.

   The Rename Disk Group dialog box (Figure A–11) is displayed.

Figure A–11: Rename Disk Group Dialog Box

3. In the Rename Disk Group dialog box:

   a. If the correct disk group name is not displayed in the Disk Group Name field, enter the disk group name or click on Browse to select the disk group.

   b. Enter the new name for the disk group.

   c. Click on Ok.

A.4.8 Moving a Disk Group to Another System

You can move a disk group (and LSM objects in that disk group) from one system to another. LSM and the Storage Administrator (server) must be running on both systems.

To move a disk group from one system to another:

1. Stop and unmount all volumes in the disk group to be moved.
2. Follow the instructions in Section A.4.4 to deport the disk group to be moved to the other system.
3. Attach all the physical disks in the disk group to the new system.
4. On the new system, follow the instructions in Section A.4.5 to import the disk group.
5. In the Object Tree, select the Storage Administrator.
6. From the System menu, choose Scan Disks to set up the newly attached disks on the system.
7. Follow the instructions in Section A.4.6 to restart and recover all volumes in the disk group on the new system.

A.4.9 Destroying a Disk Group

You can destroy a disk group permanently to remove the group from LSM control. This process reinitializes all the disks in the disk group as empty disks and places them in the free disk pool for reuse. You cannot destroy a disk group if any volumes in that disk group are in use. When a disk group is destroyed, the volumes in the disk group are removed.

Destroy a disk group only if you are sure that you no longer need the volumes and data in the disk group. Because the last disk in an existing disk group cannot be removed, destroying a disk group is a way to free the last disk in a disk group for reuse.

You cannot destroy the rootdg disk group.

To destroy a disk group:

1. In the Object Tree, select Storage Administrator.
2. From the System menu, choose Destroy Disk Group.

   The Destroy Disk Group dialog box is displayed.
3. In the Destroy Disk Group dialog box:

   a. Enter the name of the disk group to be destroyed in the Disk Group Name field or click on Browse to select the disk group.

   b. Click on Ok.

A.5 Managing Subdisks

Subdisks are created as the result of creating a volume. You cannot use the Storage Administrator to create subdisks.

The following sections describe the subdisk management tasks that you can complete using the Storage Administrator.

A.5.1 Joining Subdisks

You can join two or more subdisks to form a single, larger subdisk. Subdisks can be joined only if they belong to the same volume and occupy adjacent regions of the same disk and mirror. The joined subdisk can retain the name of one of the subdisks being joined.

For a volume with a striped plex, the subdisks must be in the same column.

To join subdisks:

1. In the Object Tree, select Volumes and in the Object Table, select the volume with the subdisks to be joined.
2. From the Volumes menu, choose Show Layout.
3. In the Volume Layout Details window, hold down the Shift key and click to select the subdisks to be combined. Subdisks must be contiguous.
4. From the Subdisks menu, choose Join.

The Join Subdisks dialog box (Figure A–12) is displayed.

Figure A–12: Join Subdisks Dialog Box

5. In the Join Subdisks dialog box:

   a. If the correct disk group name is not displayed in the Disk Group Name field, enter the name of the disk group that contains the subdisks to be joined.

   b. If the correct subdisk names are not displayed, enter the subdisk names or click on Browse to select the subdisks. Specify at least two subdisk names separated by a space.

   c. In the Target Subdisk Name field, enter the name of the new, combined subdisk.

   d. Click on Ok.

6. Optionally, view the new volume layout by selecting Update from the View menu.

7. Close the Volume Layout Details window.

A.5.2 Splitting a Subdisk

You can divide a subdisk into two or more subdisks. Once split, the new, smaller subdisks can be moved elsewhere or rejoined later. This feature is useful for reorganizing volumes or for improving performance. The original subdisk must contain a sufficient number of sectors for the specified split to work.

The name of the first new subdisk remains the same as the original subdisk. The other subdisks are automatically named by the Storage Administrator. The new, smaller subdisks occupy the same regions of the disk that the original subdisk occupied.

You cannot split a log subdisk.

To split a subdisk into multiple subdisks:

1. In the Object Tree, select Volumes and in the Object Table, select the volume with the subdisk to be split.
2. From the Volumes menu, choose Show Layout.
3. In the Volume Layout Details window, select the subdisk to be split into multiple subdisks.
4. From the Subdisks menu, choose Split.

   The Split Subdisk dialog box (Figure A–13) is displayed.

Figure A–13: Split Subdisk Dialog Box

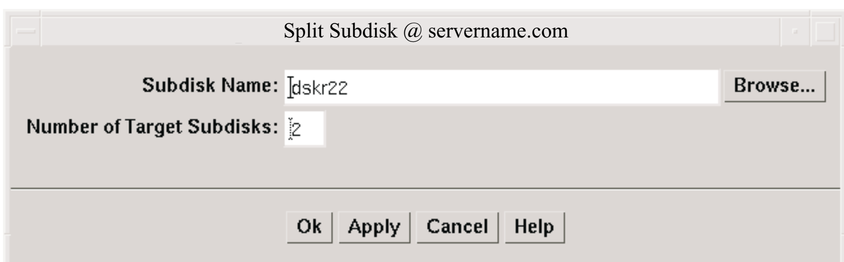

5. In the Split Subdisk dialog box:

   a. If the correct subdisk name is not displayed, enter the subdisk name or click on Browse to select the subdisk.

   b. In the Number of Target Subdisks field, enter the number of subdisks into which the subdisk should be split. A subdisk can be split into two or more subdisks.

   c. Click on Ok.

6. Optionally, view the new volume layout by selecting Update from the View menu.
7. Close the Volume Layout Details window.

A.5.3 Moving a Subdisk

You can move the data in subdisks to a different disk to improve performance. The disk space occupied by the original subdisk is returned to the free space pool.

Do not move a subdisk in a mirrored, striped, or RAID 5 volume to a disk that already contains a copy or part of that volume.

If the process of moving a subdisk leaves some unused subdisks (that is, subdisks that are not associated with a volume) on the system, you can remove the subdisk (Section A.5.4) to free the space occupied by the unused subdisks.

To move a subdisk:

1. In the Object Tree, select Volumes and in the Object Table, select the volume with the subdisk to be moved.
2. From the Volumes menu, choose Show Layout.
3. In the Volume Layout Details window, select the subdisk to be moved to another disk.
4. From the Subdisks menu, choose Move.

   The Move Subdisks dialog box (Figure A–14) is displayed.

Figure A–14: Move Subdisks Dialog Box

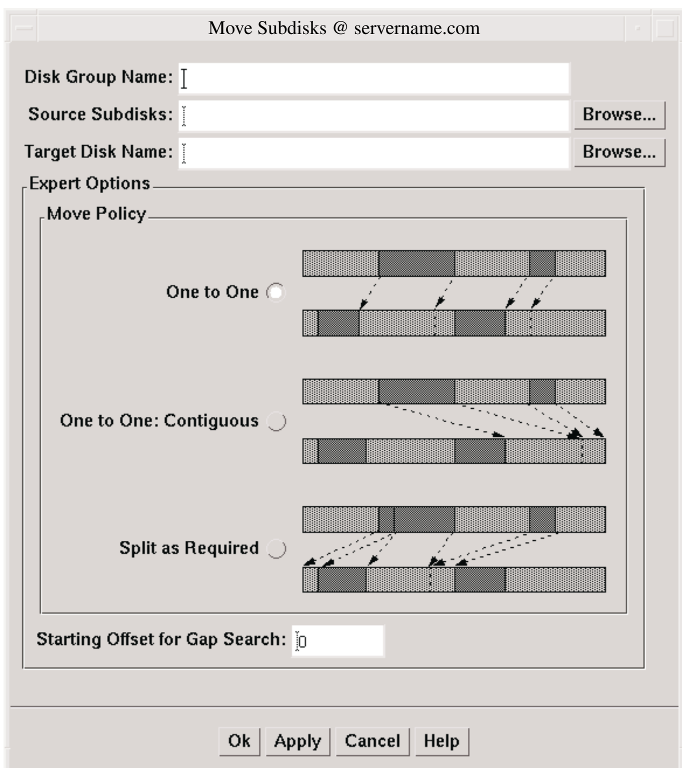

5. In the Move Subdisks dialog box:

   a. If the correct disk group name is not displayed in the Disk Group Name field, enter the name of the disk group that contains the subdisk to be moved.

b. If the correct source subdisk name is not displayed in the Source Subdisks field, enter the subdisk's name or click on Browse to select the subdisk.
c. In the Target Disk Name field, enter the name of the target disk to which the subdisk should be moved or click on Browse to select a disk.
d. Choose the Move Policy to specify whether the subdisk can be split into smaller subdisks that fit in available space(s) on the target disk. The One to One options do not split the subdisk. The Split as Required option allows the subdisk to be split if needed.
e. In the Starting Offset for Gap Search field, enter the minimum disk offset for the subdisk.

6. Click on Ok.
7. Optionally, view the new volume layout by selecting Update from the View menu.
8. Close the Volume Layout Details window.

A.5.4 Removing a Subdisk

You can remove a subdisk that is not associated with a volume. The disk space occupied by the unused subdisks is returned to the free space pool.

To remove a subdisk:

1. In the Object Tree, select Disks and in the Object Table, select the disk with the subdisk to be removed.
2. From the Disks menu, choose Properties.
3. In the Disk Properties window, click the Subdisks tab.
4. Select the subdisk to remove.
5. From the Subdisk menu, choose Remove.

   In the Remove Subdisks dialog box, click Yes to remove the subdisk.
6. Close the Disk Properties window.

A.6 Managing LSM Volumes

The following sections describe the volume management tasks that you can complete by using the Storage Administrator. Most tasks described in this section are appropriate only for UFS.

Notes

AdvFS volumes are not the same as LSM volumes. This section deals with LSM volumes. For information about AdvFS volumes, see Section A.7.1.

In a cluster, AdvFS file systems are supported in all modes; UFS file systems are supported in read-only mode.

A.6.1 Creating a Volume

You can create a volume that is less than or equal to the available free space on the LSM disks. If you do not specify the disks, the Storage Administrator uses any available space on disks in the selected disk group.

The data in a striped or concatenated volume is not protected against disk failure unless the volume is mirrored.

To create a volume:

1. From the Console menu, choose New then Volume.

   The New Volume dialog box (Figure A–15) is displayed.

Figure A–15: New Volume Dialog Box

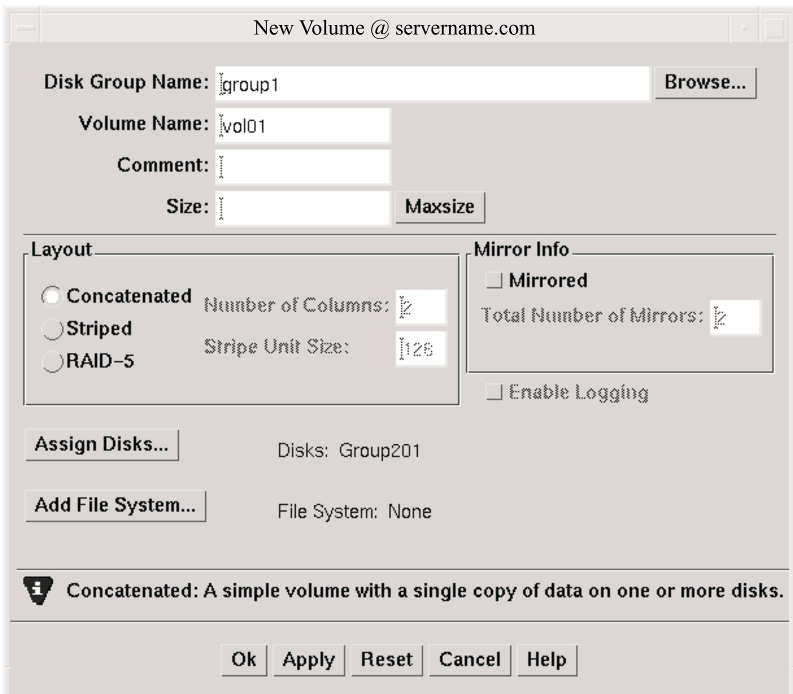

2. In the New Volume dialog box:

   a. If the correct disk group name is not displayed, enter the disk group name or click on Browse to select the disk group.

   b. Accept the default new volume name or enter a new volume name.

   c. Enter a comment if desired.

   d. Enter the volume size:

      - To specify a size unit, attach an `s` (sectors), `k` (kilobytes), `m` (megabytes), or `g` (gigabytes) to the size. The default size unit is sectors.
      - To determine the largest possible size for the volume, click Maxsize. Units are displayed in kilobytes.

   e. Choose the volume layout:

      - Concatenated
      - Striped—Enter the number of columns and stripe unit size.
      - RAID 5—Enter the number of columns and stripe unit size. This option is not available in a cluster.

f. If you have chosen a concatenated or striped volume, you can choose to mirror it.

   i. To mirror the volume, select Mirrored.

   ii. In the Total Number of Mirrors field, enter the total number of mirrors for the volume.

       Note that each plex is a mirror, so if you create a volume with one mirror, the total number of plexes is 2.

g. The Enable Logging box is selected by default when you create a mirrored or RAID 5 volume. To disable logging for the volume, deselect Enable Logging.

h. To place the volume on a specific disk, click Assign Disks.

   - Select the disk you want to use from the Space Allocation–New Volume dialog box and click on Ok.

i. To place a file system on the volume, click Add File System.

j. In the Add File System dialog box:

   i. Enter the mount point for the file system. The mount point must be an absolute pathname; that is, it must begin with root (/). If the path specified for the mount point does not exist, it is created.

   ii. Select Mount at Boot if you want the `/etc/fstab` file automatically updated and the file system mounted at system startup.

   iii. Select the file system type. If you select AdvFS, enter the Domain Name and the Fileset name. If you select UFS, you can also select Extra Options in the Newfs Details dialog box. This option is not available in a cluster environment.

   iv. Click on Mount Details to make the file system read-only or to add extra options. Click on Help in the Mount Details dialog box for more information.

   v. Click on Ok to close the Add File System dialog box.

k. Click on Ok to close the New Volume dialog box.

A.6.2 Adding a Mirror to a Volume

You can create a mirror (copy) of a volume on a disk that is not being used. When mirrored, the data in the volume is redundant. If a disk fails, the data remains available on another mirror. A volume can have multiple mirrors, but each must reside on a separate disk. Sufficient disk space must be available. You cannot mirror a RAID 5 volume.

You can use only disks in the same disk group to create a mirror. If you do not specify the disks, LSM uses available disk space to create the mirror. Adding a mirror requires resynchronization, so this task might take some time.

A volume can contain up to 32 plexes, which can be any combination of mirrors and logs.

To add one or more mirrors to an existing volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to be mirrored.
2. From the Volumes menu, choose Mirror then Add.

   The Add Mirror dialog box (Figure A–16) is displayed.

Figure A–16: Add Mirror Dialog Box

3. In the Add Mirror dialog box:

   a. If the correct volume name is not displayed in the Volume Name field, enter the volume name or click on Browse to select the volume.

   b. Click on Layout to:

      - Specify the layout for the mirror (concatenated or striped) and, if striped, the stripe unit size.
      - Add more than one mirror and supply comments.

   c. Click on Assign Disks to place the mirror on a specific disk.

   d. Click on Ok.

A.6.3 Adding a Log to a Mirrored or RAID 5 Volume

When you add a log to a mirrored volume, dirty region logging (DRL) is activated for that volume. DRL uses the log to track the regions of the volume that change due to I/O writes. If a system failure occurs, DRL uses the information in the log to recover only the portions of the volume that need recovery. This speeds up recovery time for mirrored volumes.

You can create additional DRL logs (on different disks) to mirror the DRL information.

When you add a log to a RAID 5 volume, LSM maintains in the log a copy of the data and parity being written to the volume at any given time. The RAID 5 log speeds up the resynchronization time after a system failure. If a system failure occurs, when the system starts up LSM copies the data and parity that was being written at the time of failure from the log to the appropriate areas of the RAID 5 volume.

You can create multiple RAID 5 logs (on different disks) to mirror the log information. Ideally, each RAID 5 volume should have at least two logs.

To add a log to a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to contain the log.
2. From the Volumes menu, choose Log then Add.

   The Add Log dialog box (Figure A–17) is displayed.

Figure A–17: Add Log Dialog Box

3. In the Add Log dialog box:

   a. If the correct volume name is not displayed in the Volume Name field, enter the volume name or click on Browse to select the volume.

   b. To place the log on a specific disk, enter the name of the disk in the Disk Name field or click on Browse to select a disk.

   c. Click on Ok.

A.6.4 Mounting a UFS File System on a Volume

You can mount an existing UFS file system on a volume. You can choose to have the `/etc/fstab` file automatically updated.

To mount a UFS file system on an existing volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume on which to mount the UFS file system.
2. From the Volumes menu, choose File System then Mount.

   The Mount File System dialog box (Figure A–18) is displayed.

Figure A–18: Mount File System Dialog Box

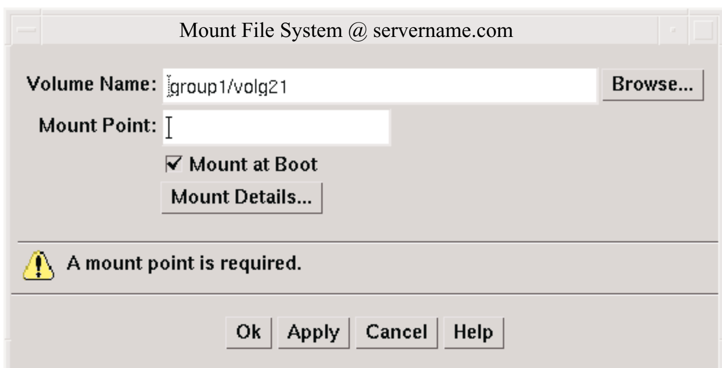

3. In the Mount File System dialog box:
   a. If the correct volume name is not displayed in the Volume Name field, enter the volume name or click on Browse to select the volume.
   b. In the Mount Point field, enter the mount point for the file system. The mount point must be an absolute pathname; that is, it must begin with root (/). If the path specified for the mount point does not exist, it is created.
   c. Select Mount at Boot if you want the `/etc/fstab` file updated and the file system mounted at system startup.
   d. Click on Mount Details to make the file system read-only or to add extra options. Click on Help in the Mount Details dialog box for more information.
   e. Click on Ok.

A.6.5 Adding a File System to a Volume

You can place a new file system on an existing volume and mount the file system. You can choose to have the `/etc/fstab` file automatically updated.

To add a file system to an existing volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to contain the file system.

2. From the Volumes menu, choose File System then New.

   The New File System dialog box (Figure A–19) is displayed.

Figure A–19: New File System Dialog Box

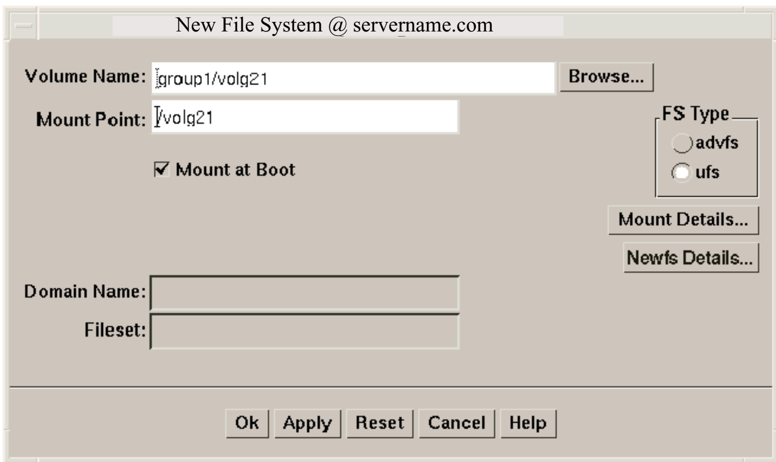

3. In the New File System dialog box:

   a. If the correct volume name is not displayed in the Volume Name field, enter the volume name or click on Browse to select the volume.

   b. In the Mount Point field, enter the mount point for the file system. The mount point must be an absolute pathname; that is, it must begin with root (/). If the path specified for the mount point does not exist, it is created.

   c. Select Mount at Boot if you want the /etc/fstab file updated automatically, and the file system mounted at system startup.

   d. Select the file system type:

      - If you select AdvFS, enter the Domain Name and the Fileset name.

        To add a volume to an existing AdvFS domain, you must have an AdvFS Advanced Utilities license. If you supply an existing domain name and do not have a license, an error message is displayed. You can create a new domain for the volume without the Advanced Utilities license.

      - If you select UFS, you can also select Extra Options in the Newfs Details dialog box. Click on Help in the Newfs Details dialog box for more information. This option is not available in a cluster.

e. Click on Mount Details to make the file system read-only or to add extra options. Click on Help in the Mount Details dialog box for more information.

f. Click on Ok.

A.6.6 Checking a UFS File System in a Volume

Checking a file system applies only to UFS file systems; AdvFS does not use the `fsck` utility. You can check the file system with or without repairing it. Checking the file system might take some time.

If you are running a cluster, UFS file systems are mounted read-only, so the Storage Administrator cannot check the file system.

To check a UFS file system on a volume:

1. In the Object Tree, select File System and in the Object Table, select the UFS file system to check.
2. From the File Systems menu, choose Check.

   The Check File System dialog box (Figure A–20) is displayed.

Figure A–20: Check File System Dialog Box

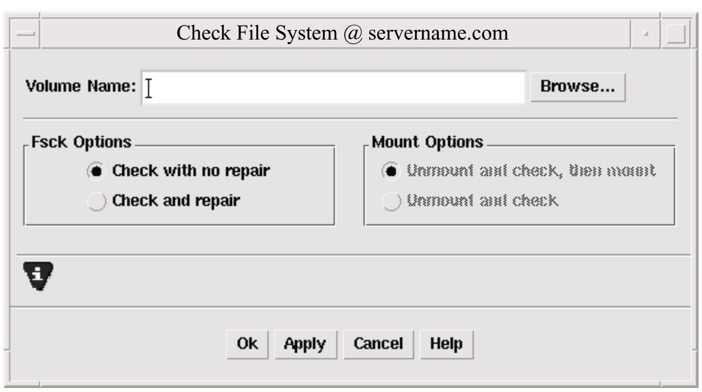

3. In the Check File System dialog box:

   a. If the correct file system name is not displayed in the File System Name field, enter the file system name or click on Browse to select the file system.

   b. Choose the Fsck option:

      - Check with no repair
      - Check and repair

   c. Choose the Mount option:

- Unmount and check, then mount
- Unmount and check

d. Click on Ok.

   When the file system check is complete, the File System Check confirmation window appears indicating that the file system is okay.

e. Click Close to close the File System Check window.

A.6.7 Unmounting a UFS File System on a Volume

You can unmount a mounted UFS file system on a volume.

To unmount a UFS file system on a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume containing the file system to be unmounted.
2. From the Volumes menu, choose File System then Unmount.

   The Unmount File System dialog box is displayed.
3. In the Unmount File System dialog box:

   a. If the correct volume name is not displayed, enter the volume name or click on Browse to select the volume.

   b. Click on Ok.

A.6.8 Mapping and Analyzing Volumes, Disks, and Subdisks

One of the most important management features of the Storage Administrator is the ability to view a map of your volumes, disks, or subdisks. Using the Volume to Disk Mapping window (Figure A–21) to view volumes or their associated disks or subdisks can save you time, especially if you have a large number of volumes and disks.

You can also use the Volume to Disk Mapping window to display information about the performance of volumes, disks, and subdisks.

The statistical values for the performance data are represented by different colors. When the analysis starts or changes, the color behind the object changes. Clicking on any of the green dots in the table highlights the path between the volume and its related disks or subdisks. Statistics can be collected only on volumes. Only disks and subdisks associated with volumes can be analyzed.

You can open Volume to Disk Mapping windows for more than one disk group; however, only one disk group can be analyzed at a time.

To open the Volume to Disk Mapping window:

1. In the Object Tree, select Disk Groups and in the Object Table, select the disk group to map.
2. From the Disk Groups men, choose Disk/Volume Map.

   The Volume to Disk Mapping window (Figure A–21) is displayed.

Figure A–21: Volume to Disk Mapping Window

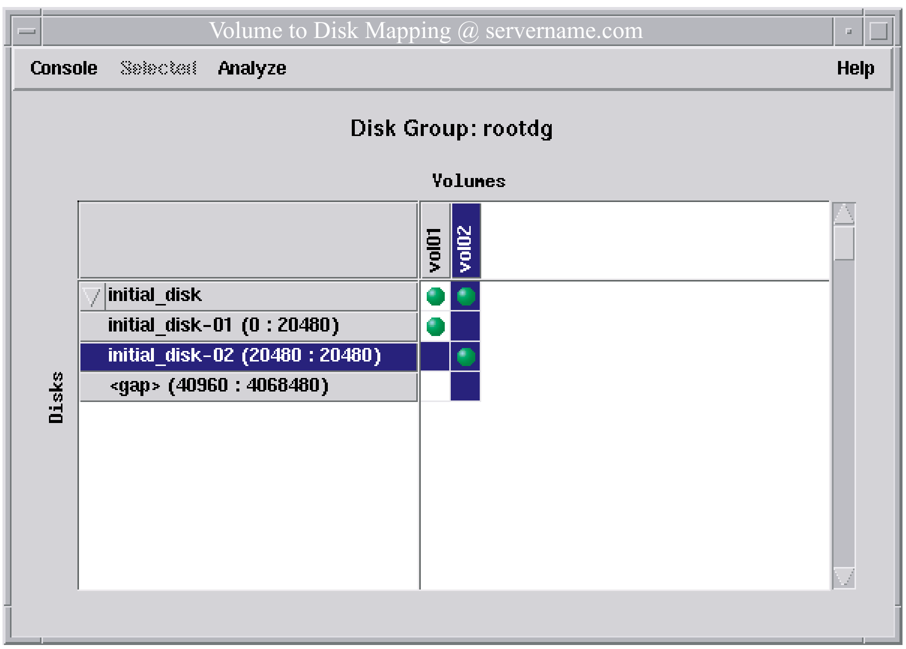

3. In the Volume to Disk Mapping window:
   - To view the associated volumes, disks, or subdisks, click on the green dot at the intersection between the disk and volume.
   - To display all the subdisks and gaps on a particular disk, click the arrow to the left of the disk name.
   - To select a disk, subdisk or volume for analysis, click on the name of that object. The background for the object changes color when selected.
   - To analyze the use of an object, select a volume, disk, or subdisk and choose a command from the Analyze menu:
     – Start Analysis—Adds the selected item or items to the list of objects being analyzed. The background color of the selected items changes to show their state during analysis.

- Stop Analysis—The selected items return to their normal state (colors disappear), while all other items being analyzed continue to display their performance characteristics.
- Analyze All—Starts analysis on all volumes and LSM disks in a view.
- Stop All—Stops analysis for all items in all views. All items return to their normal state. (The colors disappear.)
- Parameters—Opens the Analysis Parameters dialog box, which lets you set the high and low threshold values for each object under analysis.

Note

When you choose the Start Analysis and Stop Analysis command on the Analyze menu, you must have selected an object. When you choose the Analyze All command you need not select an object.

- To see the Statistics form showing numerical equivalents for the colors, right-click on the volume or disk being analyzed.
- To use the pop-up menu to perform management tasks on volumes, disks, or subdisks, right-click on the disk or volume object.
- To close the Volume to Disk Mapping window, choose Close from the Console menu.

A.6.9 Disabling a Mirror in a Volume

You can disable a mirror to temporarily detach the mirror from its volume. However, this can result in a loss of data redundancy, because the mirroring process is not occurring. A detached mirror is inaccessible for reads and writes but is still associated with its volume.

Once disabled, the mirror remains detached from its volume until you either reattach the mirror or restart the volume. When a volume is restarted, any disabled (detached) mirrors are reattached to the volume automatically.

If a volume has two mirrors and one mirror is disabled, the volume is not redundant while the mirror is disabled.

The last mirror in a volume cannot be disabled.

To disable a mirror in a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume that contains the mirror to be disabled.

2. From the Volumes menu, choose Mirror then Disable.

   The Disable Mirror dialog box (Figure A–22) is displayed.

Figure A–22: Disable Mirror Dialog Box

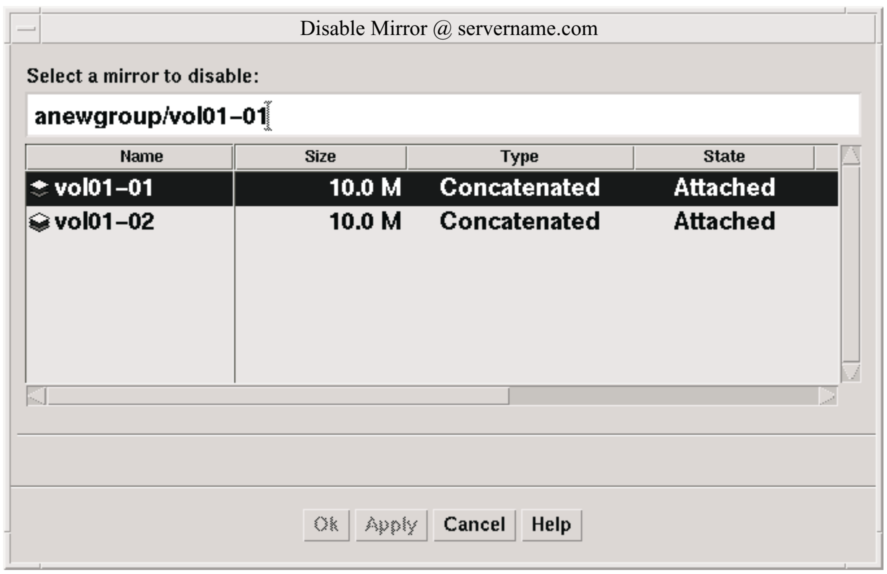

3. In the Disable Mirror dialog box:

   a. Select a mirror to disable.

      The mirror name appears in the Select a Mirror to Disable field.

   b. Click on Ok.

   In the Disable Mirror dialog box, click Yes to disable the selected mirror.

A.6.10 Repairing a Mirror in a Volume

You can repair a disabled mirror and reattach it to its volume. This process copies the data from an active mirror on the volume to the mirror being repaired and attaches the repaired mirror to the volume. Once attached, the mirror is accessible for reads and writes. The mirror is recovered so that it has the same contents as other mirrors in the volume.

Alert icons and the Alert Monitor window might indicate when a mirror needs to be repaired.

Depending on the amount of data in the volume, this task might take some time.

To repair a mirror:

1. In the Object Tree, select Volumes and in the Object Table, select the volume that contains the mirror to be repaired.
2. From the Volumes menu, choose Mirror then Repair.

   The Repair Mirror dialog box (Figure A–23) is displayed.

Figure A–23: Repair Mirror Dialog Box

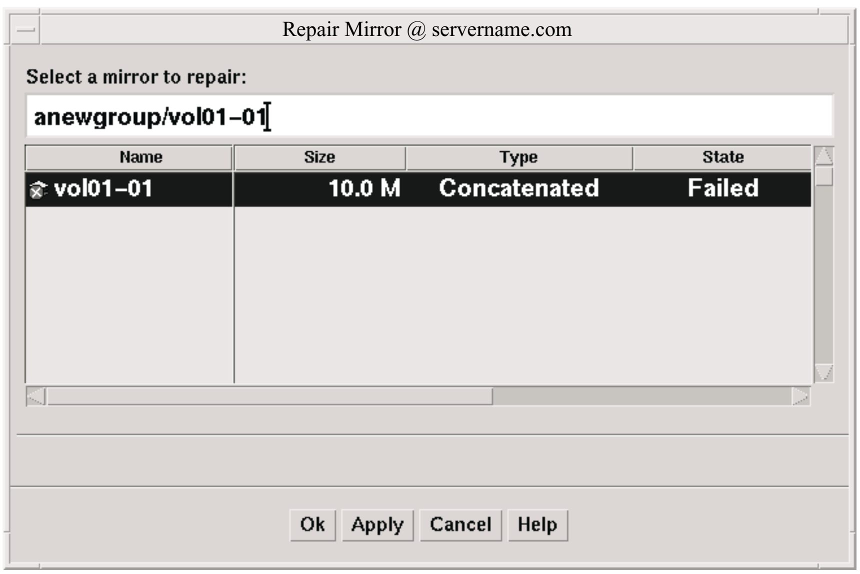

3. In the Repair Mirror dialog box:

   a. Select the mirror to be repaired.

   b. Click on Ok.

A.6.11 Creating a Snapshot of Volume Data

You can create a snapshot (temporary mirror) of a volume that you can then use to create a temporary volume for backup. You do not need to stop or take the original volume off line. In a cluster, you cannot take a snapshot of the clusterwide root.

To create or stop a volume snapshot:

1. In the Object Tree, select Volumes and in the Object Table, select a volume for which to create or stop a snapshot.
2. From the Volumes menu, choose Snapshot.

   The Volume Snapshot dialog box (Figure A–24) is displayed.

Figure A–24: Volume Snapshot Dialog Box

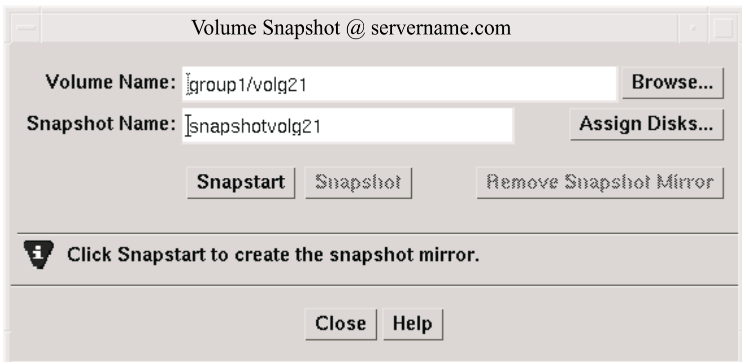

3. In the Volume Snapshot dialog box:

   a. If the correct volume name is not displayed in the Volume Name field, enter the volume name or click on Browse to select the volume.

   b. Accept the default snapshot name (volume) or enter a snapshot volume name.

   c. Click on Assign Disks to select the disks for the snapshot volume.

4. To create the snapshot, start, or stop it:

   - Click on Snapstart to create the snapshot mirror.
   - Click on Snapstop to stop the snapshot process and create a new volume to attach to the snapshot mirror so you can access it.
   - Click on Remove Snapshot Mirror to remove the temporary volume that was created.

A.6.12 Recovering a Volume

The volume recovery operations depend on the type of volume and include starting disabled volumes, resynchronizing mirrors in mirrored volumes, and resynchronizing parity in RAID 5 volumes. After successful recovery, the volume will be available for use.

Alert icons and the Alert Monitor window might indicate when a volume recovery is needed.

In some cases, recovery might not be possible. If the volume recovery fails, you can attempt to restore the volume from backup.

To recover a failed volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to be recovered.

2. From the Volumes menu, choose Recover.
3. In the Recover Volume dialog box, click Yes to recover the volume.

A.6.13 Preparing to Restore a Volume from Backup

If a volume's data is corrupted and you need to restore the volume from backup, you must prepare the volume for restoration. To restore a volume from backup, you can stop the volume, set the volume to an uninitialized state, and restart the volume (without resynchronizing the volume's mirrors). This procedure will not work for an AdvFS file domain.

If the volume contains a mounted UFS file system, you must unmount the file system before you restore the volume from backup. This task does not remount the file system.

To prepare to restore a volume from backup:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to be restored from backup.
2. From the Volumes menu, choose Prepare For Restore.

   In the Prepare Volume For Restore dialog box, click Yes to prepare the volume for restoration.

A.6.14 Removing a Mirror from a Volume

Removing a mirror from a volume breaks the link between the mirror and its volume and returns the mirror's disk space to the free space pool for reuse.

If a volume has only two mirrors and one mirror is removed, the volume is no longer redundant. This leaves the volume unprotected against disk failure. The last mirror cannot be removed from a volume, for that is equivalent to removing the volume.

To remove a mirror from a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume that contains the mirror to be removed.
2. From the Volumes menu, choose Mirror then Remove.

   The Remove Mirror dialog box is displayed.
3. In the Remove Mirror dialog box, select a mirror to remove and click Ok.

A.6.15 Removing a Log from a Volume

You can remove a DRL log or a RAID 5 log from a volume.

If you remove a volume's only log, logging (either DRL or RAID 5 logging) is no longer in effect for that volume. If logging is disabled, recovery time increases.

To remove a log from a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume that contains the RAID 5 or DRL log to be removed.
2. From the Volumes menu, choose Log then Remove.

   The Remove Log dialog box (Figure A–25) is displayed.

Figure A–25: Remove Log Dialog Box

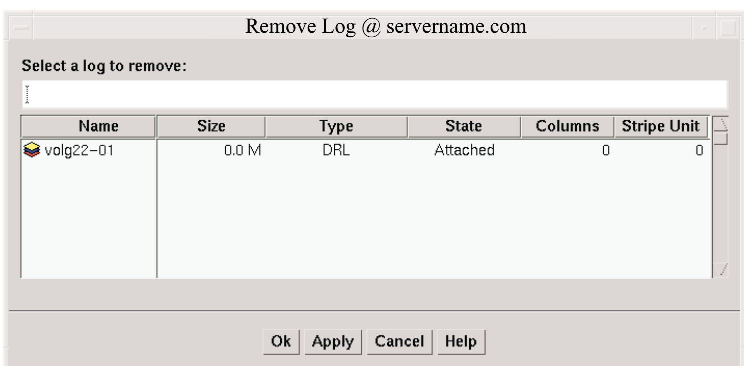

3. In the Remove Log dialog box, select the log to remove and click on Ok.

A.6.16 Removing a Volume

Removing a volume destroys all the data in that volume. Remove a volume only if you are sure that you do not need the data in the volume (or the data is backed up elsewhere). When a volume is removed, the space it occupied is returned to the free space pool.

Removing a volume that has a file system on it works only if the file system is UFS.

To remove a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to remove.
2. From the Volumes menu, choose Remove.

   The Remove Volume dialog box is displayed.
3. In the Remove Volume dialog box, click Yes to remove the volume.

A.6.17 Renaming a Volume

When you rename a volume, the new name must be unique within the disk group. If the volume has a file system, renaming the volume automatically updates the `/etc/fstab` file and allows you to specify a new mount point for the file system. You cannot rename volumes that are part of an AdvFS domain.

To rename a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to be renamed.
2. From the Volumes menu, choose Rename.

   The Rename Volume dialog box (Figure A–26) is displayed.

   Figure A–26: Rename Volume Dialog Box

   

3. In the Rename Volume dialog box:

   a. If the correct volume name is not displayed, enter the volume name or click on Browse to select the volume.

   b. Enter the new name for the volume.

   c. Click on Ok.

   If the volume is open or mounted, the Open Volumes dialog box is displayed.

4. In the Open Volumes dialog box, click Yes to rename the volume.

A.6.18 Resizing a Volume

You can increase or decrease the size of a volume, with the following exceptions:

- You cannot resize an AdvFS file domain with the Storage Administrator. If you want to resize a domain, use the AdvFS commands `addvol` and `rmvol`. See *AdvFS Administration* for more information.
- A volume containing an unmounted file system cannot be shrunk.

You can specify either the desired size or the amount of space to add to or subtract from the volume size. When a volume is shrunk, the resulting extra space is returned to the free space pool. To increase the volume size, sufficient disk space must be available. When increasing the size of a volume, LSM assigns the necessary new space from available disks.

To resize a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to be resized.
2. From the Volumes menu, choose Resize.

   The Resize Volume dialog box (Figure A–27) is displayed.

   Figure A–27: Resize Volume Dialog Box

   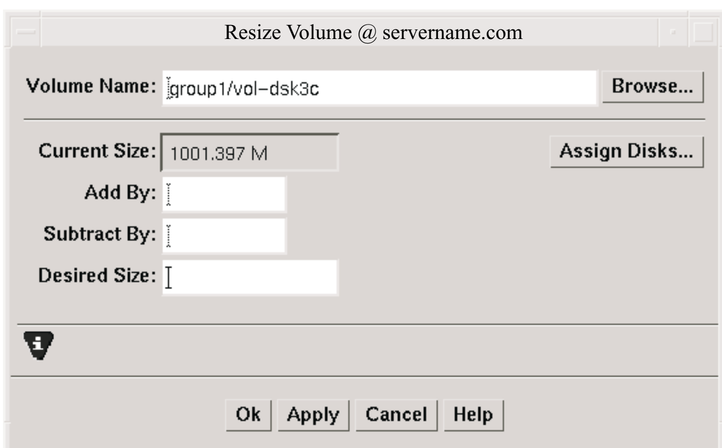

3. In the Resize Volume dialog box:

   a. If the correct volume name is not displayed, enter the volume name or click on Browse to select the volume.

   b. To use a specific disk for the additional space, click Assign Disks and select the disk you want to use from the Space Allocation–Resize dialog box.

   c. Specify one of the following:

      - To increase the volume size by a specific amount of space, use the Add By field to specify how much space to add.
      - To decrease the volume size by a specific amount of space, use the Subtract By field to specify how much space to remove.
      - To specify the new volume size, enter the size in the Desired Size field.

To specify a size unit, attach an `s` (sectors), `k` (kilobytes), `m` (megabytes), or `g` (gigabytes) to the size. The default unit is sectors.

4. Click on Ok.

A.6.19 Restarting a Volume

Under normal circumstances, volumes are automatically started when the system restarts. You can restart a volume that you stopped manually or that was stopped in some other manner. If you cannot start a volume, the volume remains unusable. If the volume contains an AdvFS file domain, you cannot start it using the following procedure.

If you are not running a cluster, restarting a RAID 5 volume enables the volume and resynchronizes parity, if necessary. Restarting a mirrored volume enables the volume and resynchronizes the mirrors to ensure that they are consistent. When a volume is successfully restarted, the volume is again available for use.

To restart a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the stopped volume to be started.
2. From the Volumes menu, choose Start.

   The Start Volume dialog box is displayed.
3. In the Start Volume dialog box:

   a. If the correct volume name is not displayed, enter the volume name or click on Browse to select the volume.
   b. Click on Ok.

      The volume's state in the Object Table changes to Started.

A.6.20 Stopping a Volume

When you stop a volume, it is not available for use until you restart it. You cannot stop a volume if it is in use or it has a mounted file system. If the volume contains an AdvFS file domain, you cannot stop it using the following procedure.

To stop a volume:

1. In the Object Tree, select Volumes and in the Object Table, select the volume to be stopped.
2. From the Volumes menu, choose Stop.

The Stop Volume dialog box is displayed.

3. In the Stop Volume dialog box:

    a. If the correct volume name is not displayed, enter the volume name or click on Browse to select the volume.

    b. Click on Ok.

       The volume's state in the Object Table changes to Stopped.

A.7 Performing AdvFS Operations

The following sections describe Advanced File System (AdvFS) tasks that you can complete using the Storage Administrator.

A.7.1 Creating an AdvFS Domain

You can create an AdvFS domain to hold all disks or partitions on your system or create several domains to hold disks specific to particular applications. When you create a domain, you must specify one volume. An AdvFS volume can be created from a complete disk, a disk partition, or an LSM volume.

Existing data on the volume you assign to a new domain is destroyed when the domain is created.

Figure A–28 shows the main Storage Administrator main window with the AdvFS Domains entry fully expanded to include Filesets and Volumes.

Figure A–28: Storage Administrator Main Window

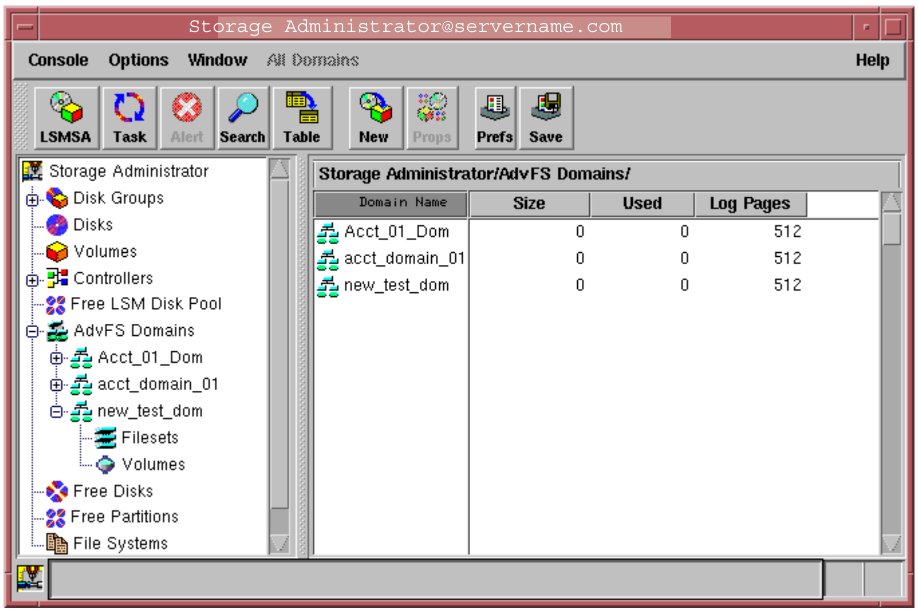

To create a domain:

1. In the Object Tree, select AdvFS Domains and from the All Domains menu choose New Domain.

   The New Domain dialog box (Figure A–29) is displayed.

Figure A–29: New Domain Dialog Box

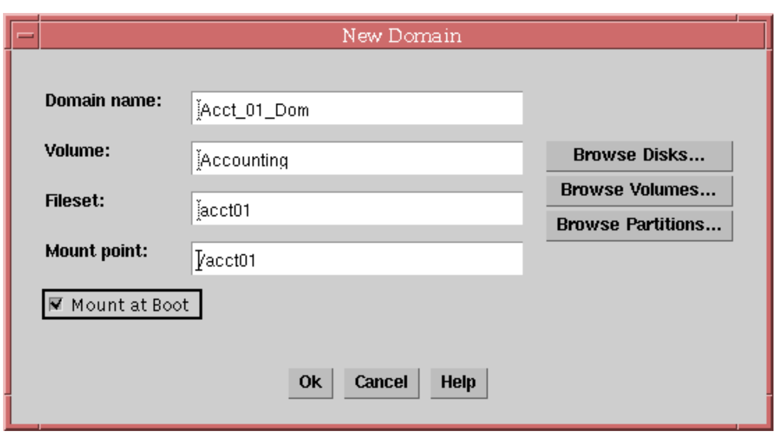

2. In the New Domain dialog box:

   a. Enter the domain name. Characters invalid for domain names are white-space characters (tab, new line, space and so on) and punctuation marks and symbols (/ # : * ?).

   b. Enter the volume name used to create the new domain, or click on Browse Disks, Browse Volumes, or Browse Partitions to select one.

   c. Enter the fileset name. Characters invalid for fileset names are white-space characters (tab, new line, space and so on) and punctuation marks and symbols (/ # : * ?).

   d. Enter the mount point for the fileset. The mount point must be an absolute pathname; that is, it must begin with root (/).

   e. Select Mount at Boot if you want the fileset to be automatically mounted at system startup.

   f. Click on Ok.

A.7.2 Adding a Volume to an AdvFS Domain

An AdvFS volume can be a single disk partition, an entire disk, or an LSM volume.

To see volumes already contained in a domain, open AdvFS Domains in the Main Window by selecting the plus sign, select the domain you want to work on, then select the volume icon.

If there are volumes within the AdvFS domain, they are displayed in the Object Table.

To add a volume to an AdvFS domain:

1. In the Object Tree, select AdvFS Domains.
2. Select the domain to contain the volume.
3. From the Domains menu, choose Add volume.

   The Add Volume dialog box (Figure A–30) is displayed.

Figure A–30: Add Volume Dialog Box

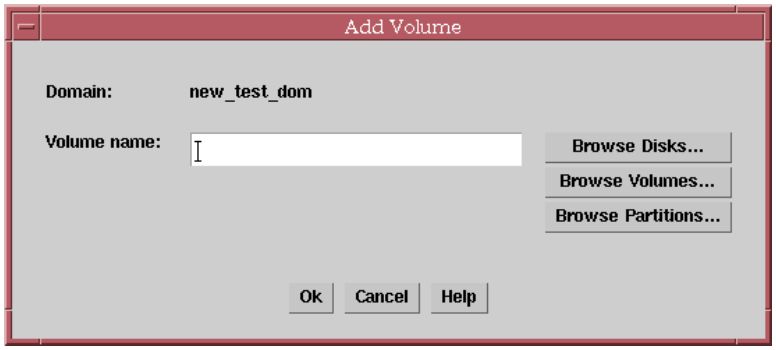

4. In the Add Volume dialog box:
   - Enter the name of the storage device to be used to create the volume and click on Ok.
   - Click Browse Disks, Browse Volumes, or Browse Partitions to view a list of possible storage devices that can be used to create the volume.
5. Enter the new name for the volume to be added to the domain.
6. Click on Ok.

A.7.3 Adding a Fileset to an AdvFS Domain

You can add a fileset to an existing AdvFS domain. Fileset names must be unique within a domain. You can choose to have the /etc/fstab file updated automatically.

You can limit the number of files in a fileset by assigning quotas.

To add a fileset to an existing AdvFS domain:

1. In the Object Tree, select AdvFS domains and in the Object Table select the domain to which to add the fileset.
2. From the Domains menu, choose Add Fileset.

   The Add Fileset dialog box (Figure A–31) is displayed.

Figure A–31: Add Fileset Dialog Box

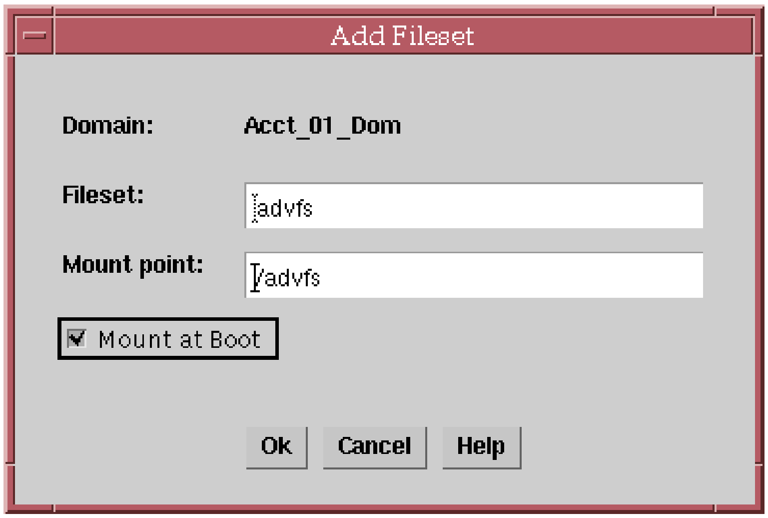

3. In the Add Fileset dialog box:

   a. In the Fileset field, enter the name of the fileset to add to the domain.

   b. In the Mount Point field, enter the mount point for the fileset. The mount point must be an absolute pathname; that is, it must begin with `/`.

   c. Select the Mount at Boot option if you want to update the `/etc/fstab` file and automatically mount the fileset at system startup.

   d. Click on Ok.

A.7.4 Defragmenting an AdvFS Domain

Before you can defragment an AdvFS domain, all filesets in the domain must be mounted. If you try to defragment an active domain that includes unmounted filesets, Storage Administrator displays an error message.

Defragment an AdvFS domain only when it is needed for system performance.

To defragment an AdvFS domain:

1. In the Object Tree, select AdvFS Domains and in the Object Table, select a domain to defragment.

2. From the Domains menu, choose Defragment.

   The Defragment Domain dialog box (Figure A–32) is displayed.

Figure A–32: Defragment Domain Dialog Box

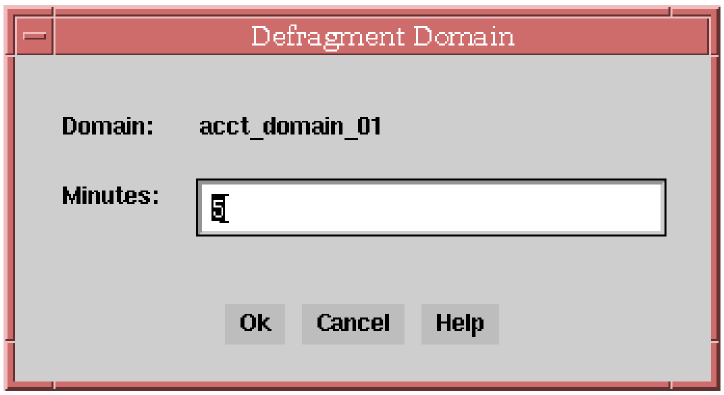

3. Enter the number of minutes for the defragment utility to run.

   If the utility is still performing an operation when the specified time has elapsed, the procedure continues until the operation is complete.

4. Click on Ok.

A.7.5 Cloning an AdvFS Fileset

At any time, only one clone can be created for an AdvFS fileset. A RAID 5 volume cannot be cloned.

To clone an AdvFS fileset:

1. In the Object Tree, click the plus sign to open the AdvFS Domains, then click the plus sign to open the domain containing the fileset to clone.
2. In the Object Tree, select Filesets.
3. In the Object Table, select the fileset you want to clone.
4. From the Filesets menu, choose Clone.

   The Clone Fileset dialog box (Figure A–33) is displayed.

Figure A–33: Clone Fileset Dialog Box

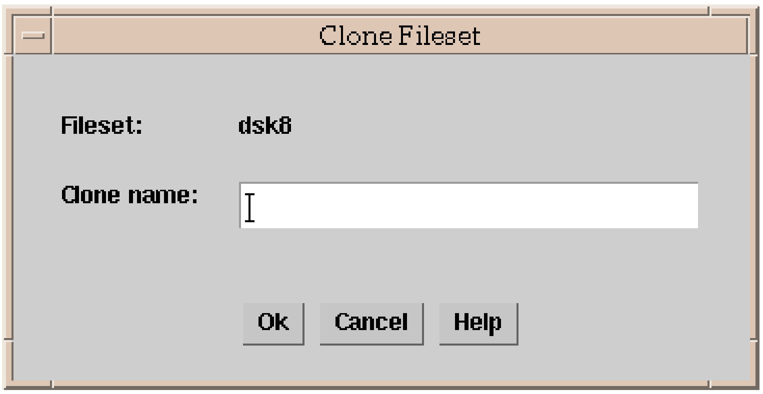

5. Enter a name for the cloned fileset.
6. Click on Ok.

A.7.6 Changing Fileset Quotas

You can change fileset quota limits. When a quota limit is reached, no more disk space allocations or file creations that would exceed the limit are allowed.

The quota numbers for block limits are displayed in kilobytes. Note that fileset quotas are not kept for the root user.

To change a fileset quota:

1. In the Object Tree, click the plus sign to open the AdvFS Domains, then click the plus sign to open the domain containing the fileset with limits to change.
2. In the Object Tree, select Filesets.
3. Select the fileset with limits to change.
4. From the Filesets menu, select Quotas.

   The Fileset Quotas dialog box (Figure A–34) is displayed.

Figure A–34: Fileset Quotas Dialog Box

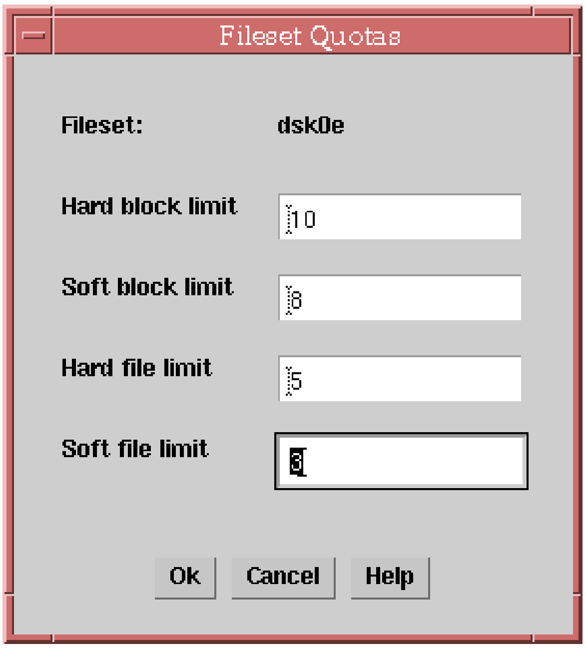

5. In the Fileset Quotas dialog box:
   - Enter the Hard and Soft Block limits in whole kilobyte increments
   - Enter the Hard and Soft file limits in whole number increments
6. Click on Ok.

A.7.7 Unmounting an AdvFS Fileset

You must unmount an AdvFS fileset to rename the fileset or to resize the underlying LSM volume.

To unmount a fileset:

1. In the Object Tree, click the plus sign to open the AdvFS Domains, then click the plus sign to open the domain containing the fileset to unmount.
2. In the Object Tree, select Filesets.
3. In the Object Table, select the fileset to unmount.
4. From the Filesets menu, choose Unmount.
5. Verify that the name of the fileset is correct.
6. Click on Yes.

A.7.8 Renaming an AdvFS Fileset

You can rename an AdvFS fileset. You must unmount the fileset before you can rename it. When you rename a fileset, the `/etc/fstab` file is updated.

To rename an AdvFS fileset:

1. In the Object Tree, click the plus sign to open the AdvFS Domains, then click the plus sign to open the domain containing the fileset to rename.
2. In the Object Tree, select Filesets.
3. In the Object Table, select the fileset you want to rename.

   The Rename Fileset dialog box (Figure A–35) is displayed.

Figure A–35: Rename Fileset Dialog Box

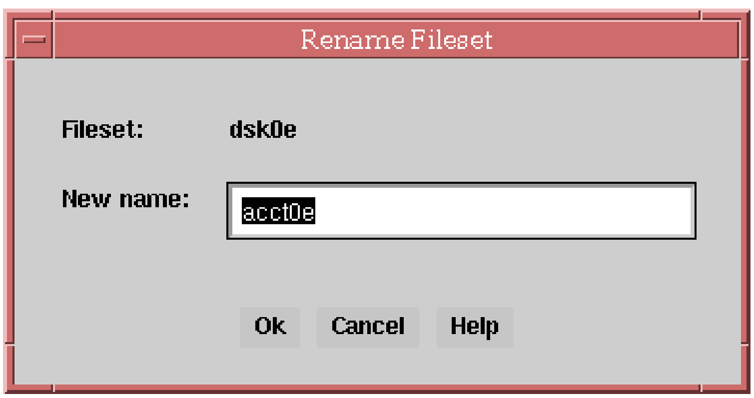

4. Enter the new name for the fileset.
5. Click on Ok.

A.7.9 Mounting an AdvFS Fileset

You can mount an AdvFS fileset that you unmounted to rename or to resize the underlying LSM volume.

To mount an AdvFS fileset:

1. In the Object Tree, click the plus sign to open the AdvFS Domains, then click the plus sign to open the domain containing the fileset to mount.
2. In the Object Tree, select Filesets.
3. In the Object Table, select the fileset to mount.
4. From the Filesets menu, choose Mount.

   The Mount Fileset dialog box (Figure A–36) is displayed.

Figure A–36: Mount Fileset Dialog Box

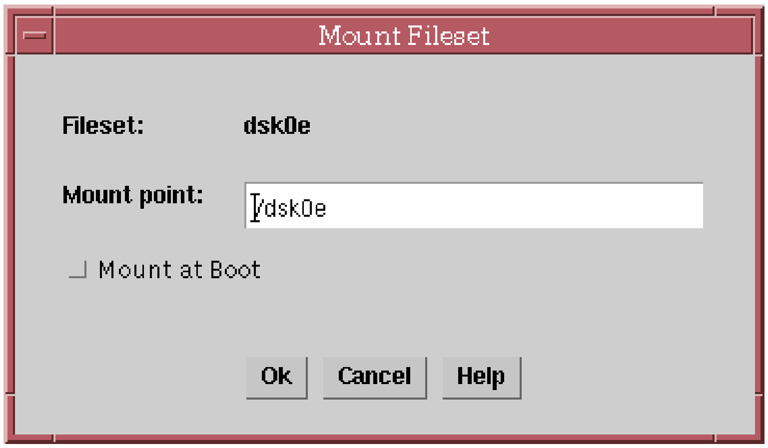

5. In the Mount Fileset dialog box:
   - Select Mount at Boot if you want to automatically update the `/etc/fstab` file and mount the fileset at system startup.
   - Enter a mount point for the fileset. The mount point must be an absolute pathname; that is, it must begin with `/`.
6. Click on Ok.

A.7.10 Removing an AdvFS Fileset

To remove an AdvFS fileset:

1. In the Object Tree, click the plus sign to open the AdvFS Domains, then click the plus sign to open the domain containing the fileset to remove.
2. In the Object Tree, select Filesets.
3. In the Object Table, select the fileset you want to remove.
4. From the Filesets menu, choose Remove.

   The Remove Fileset window is displayed.
5. Click on Yes to remove the indicated fileset.

B

Understanding and Customizing the Storage Administrator GUI

This appendix describes how to track Storage Administrator activities, how to use the Storage Administrator dialog box, how to view objects and object properties, and how to use shortcuts to efficiently perform operations with the Storage Administrator GUI. See Appendix A for more information on using the Storage Administrator to complete a particular task for an LSM or AdvFS object.

B.1 Tracking Storage Administrator Activities

Three log files keep track of the Storage Administrator:

- A command log tracks Storage Administrator tasks
- An access log tracks Storage Administrator logins
- A server log collects LSM and startup information

  The same server log collects LSM and AdvFS information and error messages.

By default, a log maintenance shell script called `/usr/lib/java/applications/lsmsa/logMaintenance` runs once a week to save and compress each log file. Compressed files are saved in a file called *`logfilename`*`.gz.X`, where `X` is the version number. Each week the previous week's saved file suffix is increased by one and a new *`logfilename`*`.gz.1` is created. Files are saved for ten weeks. You can change the number of files to save by editing the root `crontab` file.

B.1.1 Command Log File

The command log file contains a description of each Storage Administrator task and information such as the user who performed the task, the task status, the start and finish times, and the commands used to perform the task. For failed tasks, the command log includes relevant error messages. By default, the command log is `/var/lsmsa/logs/command`.

The following output shows a sample command log entry for a successful volume creation:

```
Create Volume
Description: Create Volume
User: root
Started: Fri Mar 09 12:07:22 PDT 2001
Finished: Fri Mar 09 12:07:24 PDT 2001

State: Successful
Executed Commands:

/usr/sbin/volassist
-g rootdg make vol04 4m layout=striped stripeunit=128 ncolumn=2
```

The following output shows a sample command log entry for a failed volume creation:

```
Create Volume FAILED!
Description: Create
VolumeUser: root
Started: Fri Mar 09 12:07:50 PDT 2001
Finished: Fri Mar 09 12:07:51 PDT 2001
State: Failed

Executed Commands:
/usr/sbin/volassist
-g rootdg   make vol05 8g layout=striped stripeunit=12 ncolumn=2

Failed Command: /usr/sbin/volassist
-g rootdg   make vol05 8g  layout=striped stripeunit=128 ncolumn=2

Error Message: lsmsa:volassist: ERROR: Cannot allocate space
for  16777216 block volume
```

B.1.2 Access Log File

You can monitor access to the Storage Administrator by reviewing the contents of the access log file. By default, the access log file is `/var/lsmsa/logs/access`.

The following output shows a sample access log file entry:

```
Mon Apr 02 12:07:22 PDT 2001: user rssn login succeeded

Mon Apr 02 12:22:24 PDT 2001; user jehg login failed with error
*User password invalid*
```

Entries for failed access might be logged multiple times due to a security requirement.

B.1.3 Server Log File

The server log file tracks LSM startup information and server errors. By default, the server log file is `/var/lsmsa/logs/server.log`.

The following output shows sample server log file entries:

```
Starting Compaq Storage Administrator RMI Registry
Starting Compaq Storage Administrator Command Server
Starting Compaq Storage Administrator Server
Fri Mar 16 11:22:21 PST 2001
```

```
security enabled
rebinding ....
rebound

 //servername:2410/vrts.remote.vrtsServer
```

B.2 Working with Dialog Boxes

The Storage Administrator displays dialog boxes in which you provide information as shown in Figure B–1. Dialog boxes can contain selectable buttons or fields in which you enter information. Some dialog box fields contain default values that you can change. Items that are not applicable are grayed out.

Figure B–1: Typical Storage Administrator Dialog Box

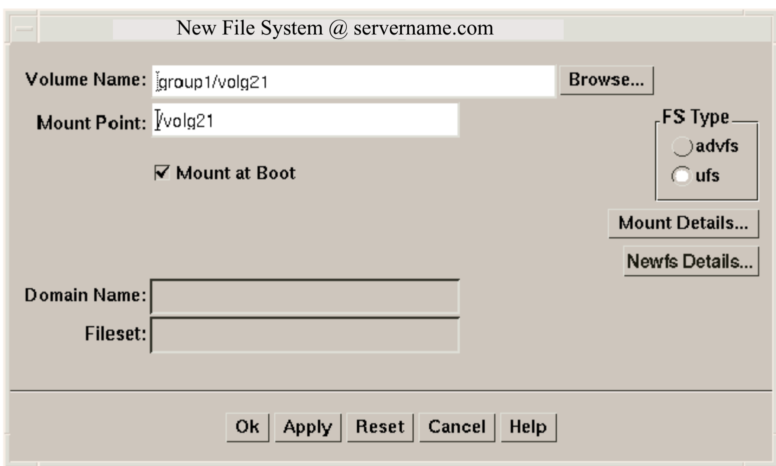

To use a dialog box, select the appropriate items or enter the appropriate information in a field, then click on one of the following buttons to initiate or cancel the task:

- Ok — Performs the current task and closes the dialog box.
- Apply — Performs the current task and continues to display the dialog box.
- Cancel — Closes the dialog box and cancels the current task. If you have already chosen Apply, use this button to close the dialog box. Doing so does not cancel the Apply request.
- Reset — Clears the information in dialog box fields.
- Help — Displays the Help menu.

B.2.1 Specifying Objects in Dialog Boxes

Most Storage Administrator dialog boxes contain one or more object name fields. If you select an object before you select the task, the resulting dialog box usually includes the selected object name. If the object name field is empty, you can specify an object using one the following methods:

- Enter the object name.

  In some cases, you can specify multiple objects (separated by a space) in a single field.

- Click on Browse next to the object name field, and then select the object from the resulting browse dialog box. Most browse dialog boxes display an Object Tree and Object Table. To select an object in a browse dialog box, click on an object group in the Object Tree, then click on the object in the Object Table.

Note

When you select an object for an action, the object name appears in the dialog box. This does not mean that the action you have chosen for that object is valid. If you try to complete an invalid operation on an object, an error message is displayed.

B.2.2 Specifying Object Sizes in Dialog Boxes

The following table shows the object sizes that you can enter to specify for an input field or a display size:

For	Enter
Sectors	s
Kilobytes	k
Megabytes	m
Gigabytes	g

By default, sectors are used for input fields if you do not specify an input size or if you did not change the default value by customizing the GUI as described in Section B.3.4. By default, sizes are displayed in kilobytes unless you specify otherwise.

B.3 Viewing Objects and Object Properties

There are several windows and dialog boxes that you can use to display information about and perform LSM operations.

B.3.1 Main Window

The Object Tree and the Object Table track your LSM configuration. The Storage Administrator constantly monitors objects on the system and makes appropriate changes to the displays. You can view objects in the Object Tree and Object Table in the following ways:

- Click on the plus sign (+) or minus sign (–) next to an object to expand or collapse its hierarchy.
- Click on the object type in the Object Tree. All objects that belong to the selected object appear in the Object Table.

  For example, to display all volumes in the rootdg disk group, expand the Disk Groups node (by clicking on the plus sign), expand the rootdg node, and click on the Volumes group under rootdg. Only volumes in the rootdg disk group appear in the Object Table.
- To display the components of an object in the Object Table, double-click on the object. All objects that belong to that object appear in the Object Table. If the object does not contain other objects, the Properties dialog box appears.

  For example, to display the volumes in a disk group listed in the Object Table, double-click on the disk group name, then double-click on Volumes. All volumes in the disk group appear in the Object Table.
- It might be useful to view a copy of the Object Table to look at different objects; for example, disks and volumes. From the Window menu choose Copy Object Table. A dynamic copy of the Object Table appears in a separate window. The contents of the new Object Table window continue to update.

B.3.2 Viewing Volume Layout Details

To display the layout details for a volume, highlight the volume in the Main window Object Table and choose Show Layout from the Volumes menu.

The Volume Layout Details window displays a graphical view of the selected volume's layout, components, and properties, as shown in Figure B–2.

Figure B–2: Volume Layout Details Window

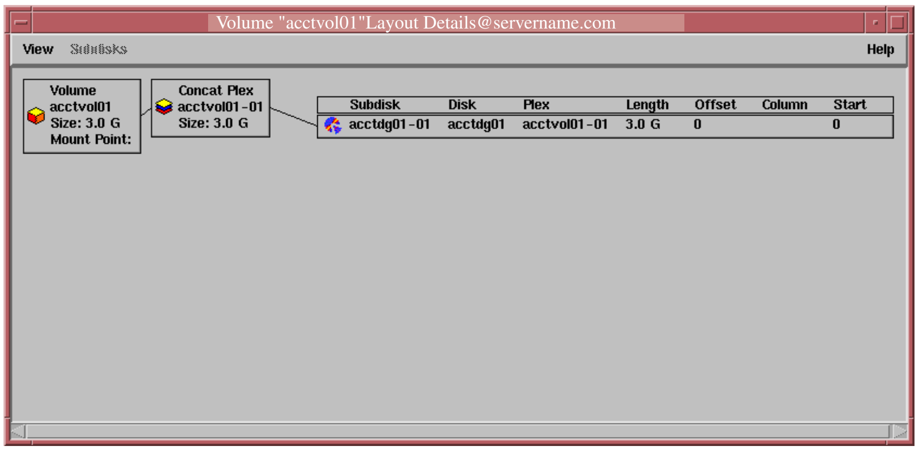

You can select objects or perform tasks on objects in the window. The Volume Layout Details window is not dynamic, so the objects displayed in the window are not automatically updated when the volume properties change. The View menu changes the display of the Volume Layout Details window.

- To refresh the display, choose Update from the View menu.
- To change the volume displayed, choose Open from the View menu and specify another volume in the Open Volume dialog box.
- To hide the detailed information within each object, choose Compress Display from the View menu. Click on an object to show its details in the compressed display.
- To highlight objects that are related to or part of a specific object, choose Projection on Selection from the View menu, then click on an object.
- To highlight any subdisks on the same disk as a specific subdisk, choose Subdisk Projection from the View menu, then click on a subdisk.

Right-click on an object to display its context-sensitive pop-up menu.

B.3.3 Object Properties Dialog Boxes

To view the properties of an object, click on the object in the Object Table and then choose Properties from the object's menu. If the object contains no other objects, double-click on the object to display its Properties dialog box.

The Object Properties dialog box displays detailed information specific to the selected object as shown in Figure B–3.

Figure B–3: Volume Properties Dialog Box

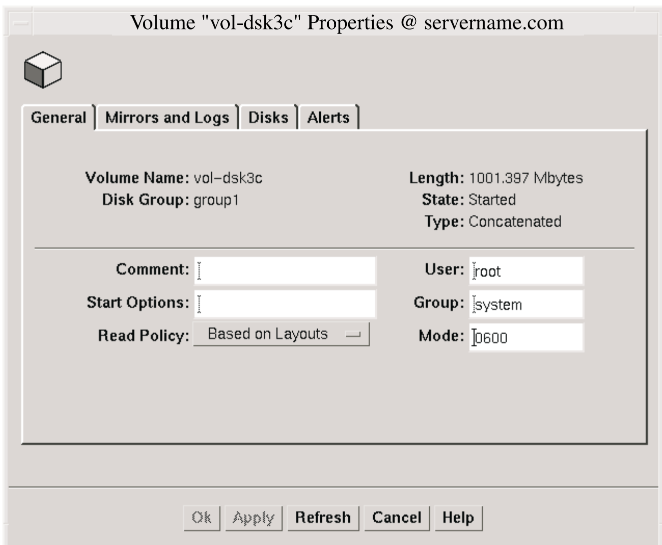

You can change some properties through this box. A set of tabbed pages provides information about the object and related objects. The tab labels and page contents vary, depending on the type of object selected. Click on Help for a detailed description of the Properties dialog box fields.

- To change items in the Properties dialog box, make the changes, then click on Ok. This changes the settings for all properties tabs in the Properties dialog box.
- To update the contents of the Properties dialog box to reflect current properties for the object, click on Refresh. If you select a different object while a Properties dialog box is open, the contents of the dialog box does not change to reflect the new object selected. You must choose Properties again and open another dialog box.

B.3.4 User Preferences Dialog Box

You can change the way items appear in the Storage Administrator Main window and other windows. The Preferences dialog box (Figure B–4) contains a set of tabbed pages that display preference options for a particular aspect of Storage Administrator. You can customize settings for a single Storage Administrator session or save the settings for future sessions.

To display the Preferences dialog box, click on Prefs on the toolbar.

Figure B–4: Preferences Dialog Box

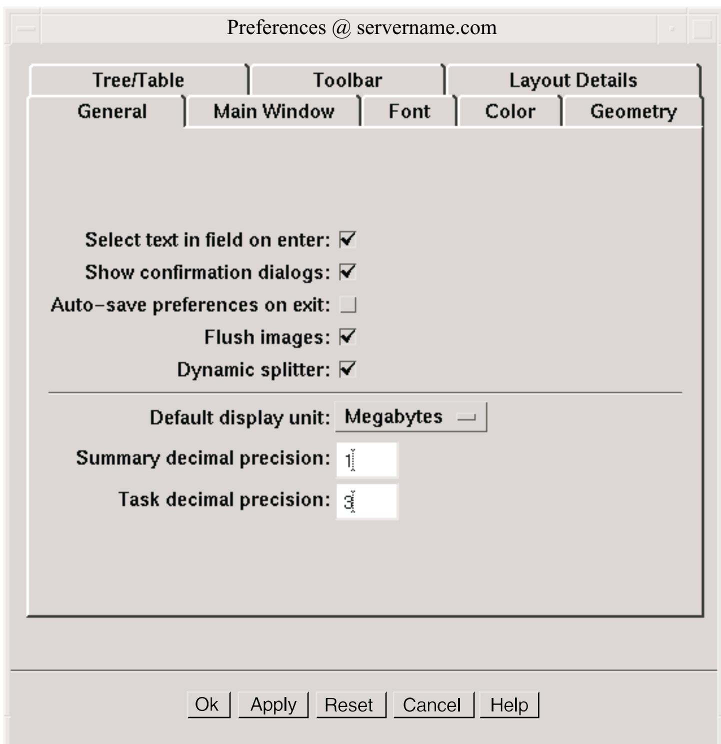

To change preference settings, make the appropriate selections in the dialog box, then click on Ok. This changes the settings for all tabs in the Preferences dialog box. To reset the values for all preferences to the previous settings, click on Reset before you click on Ok.

When you change preference settings, an asterisk appears on the tab for the page that contains changes. The asterisk disappears when you click on Ok, Apply, or Reset. When you click on Apply or Reset, an asterisk also appears in the Help bar status area in the Storage Administrator Main Window.

Unless you save your preferences, changes apply only to the current session. To save your settings, do one of the following:

- From the Options menu, choose Save Preferences
- Click on Save in the toolbar
- Click the asterisk in the Help bar status area

To reload your previously saved preferences, choose Load Preferences from the Options menu.

The Storage Administrator saves user preferences in the *user's_home_directory*/`.lsmsa/SApreference.prf` file on the system where the client is running. If the autosave preference is set, the Storage Administrator saves all preference settings when you exit the Storage Administrator session.

B.3.4.1 General Preferences

The General tab sets the preferences for:

- Select Text in Field on Enter

  Sets user input to replace mode. This highlights existing text in a field and replaces that text with the new text.

- Show Confirmation Dialogs

  Shows or hides confirmation dialog boxes for tasks that might have serious consequences (such as data loss). Confirmation dialogs require you to confirm that a task be performed. Confirmation dialogs typically appear for tasks that remove objects. If you hide confirmation dialogs, most tasks are performed immediately and without any confirmation.

- Auto-Save Preferences on Exit

  Saves all current user preferences when you exit the Storage Administrator.

- Flush Images

  Draws images slightly slower than usual to prevent the X server from growing. This is recommended if you plan to run the Storage Administrator for long periods of time.

- Dynamic Splitter

  Refreshes the contents of the window panes while the splitter is being moved to resize the panes.

- Default Display Unit

  Sets the default size unit for areas that display object sizes. If Best Choice is set, the Storage Administrator uses an appropriate size unit.

- Summary Decimal Precision

  Sets the decimal point precision for object sizes displayed in the Object Table and other areas that display summaries.

- Task Decimal Precision

  Sets the decimal point precision for object sizes displayed in task-related dialog boxes and areas that display numerical information.

B.3.4.2 Main Window Preferences

The Main Window tab sets the preferences for:

- Show Status Bar

  Shows or hides the status bar (at the bottom of the Main window). The Status bar displays alert icons when failures or errors occur.

- Show Command Launcher

  Shows or hides the Command Launcher. The Command Launcher displays a list of selectable tasks. You can show or hide the Command Launcher by checking or unchecking the Command Launcher box on the Window menu of the Main window.

- Dock Command Launcher

  Attaches or detaches the Command Launcher and the Main window.

- Docked Command Launcher Height

  Sets the height of the Command Launcher portion of the Main window.

B.3.4.3 Font Preferences

The Font tab sets the font size, family, and style for:

- User Font

  Sets the font for user input and objects displayed in the Object Tree and Object Table.

- System Font

  Sets the font for the Storage Administrator labels, menus, and buttons.

- Object Table Heading Font

  Sets the font for Object Table headings.

- Object Table Heading Highlight Font

  Sets the font for the highlighted Object Table headings for sorting purposes.

- Toolbar Font

  Sets the font for the toolbar buttons.

- Graphical Display Font

  Sets the font for objects in the Volume Layout Details window.

B.3.4.4 Color Preferences

The Color tab sets color preferences. Change colors by clicking on a color in the color wheel or by sliding the Red, Green, Blue, and Brightness sliders.

Colors can be set for:

- Background Color

  Sets the background color for all the Storage Administrator windows.
- Foreground Color

  Sets the color for foreground text in the Storage Administrator windows.
- Tree/Table Color

  Sets the background color for the Object Tree and Object Table.
- Connecting Line Color

  Sets the color for the lines that connect items in the Object Tree.
- Selection Color

  Sets the color for selected items.
- Selection Foreground Color

  Sets the color for foreground text in selected items.
- Link Color

  Sets the color for links (such as the links to tasks in the Command Launcher).
- Projection Color

  Sets the color for the lines that show object relationships in the Volume Layout Details window.

B.3.4.5 Geometry Preferences

The Geometry tab sets the width and height (in pixels) for:

- Main window
- Object Search window
- Alert Monitor window
- Task Request Monitor window
- Volume Layout Details window
- Command Launcher window
- Object Table Copy window

If you resize one of these windows with the splitters, the new size is displayed in the Geometry preference for that window the next time you open the Preferences dialog box.

B.3.4.6 Tree/Table Preferences

The Tree/Table tab sets Object Tree and Object Table preferences for:

- Display Full Path

  Displays path information in the Object Tree and Object Table.

- Auto Scroll Table

  When an object is added or changed, scrolls through the objects until the new or changed object is visible in the Object Table.

- Splitter Position

  Moves the splitter to adjust the relative sizes of the Object Tree and Object Table panes.

- Selector Tree/Table Width

  Sets the width (in pixels) of the Object Tree and Object Table for Browse dialog boxes that contain an Object Tree and an Object Table.

- Selector Table Width

  Sets the width (in pixels) of the Object Table for Browse dialog boxes that contain only an Object Table.

- Visible Selector Rows

  Sets the number of rows displayed in the Object Tree and Object Table in Browse dialog boxes.

B.3.4.7 Toolbar Preferences

The Toolbar tab sets preferences for:

- Show Toolbar

  Shows or hides the toolbar.

- Position

  Places the docked toolbar at the top, bottom, or side of the Main window.

- Presentation

  Displays graphics, labels, or both on the buttons in the toolbar.

B.3.4.8 Layout Details Preferences

The Layout Details tab sets Volume Layout Details window preferences for:

- Compress Display

  Compresses the graphical display of objects so that details are hidden.

- Projection on Selection

When an object is selected, highlights objects that are related to or part of that object.

- Subdisk Projection

  When a subdisk is selected, highlights other subdisks on the same disk.

B.3.5 Alert Monitor Window

The Alert Monitor window (Figure B–5) displays information about failed objects or objects that experienced other errors. Each object is displayed with a description of the failure or error. When an object fails and an alert occurs, the Alert icon appears on the Status bar of the Main window and also overlays the object's icon in the Object Table.

Figure B–5: Alert Monitor Window

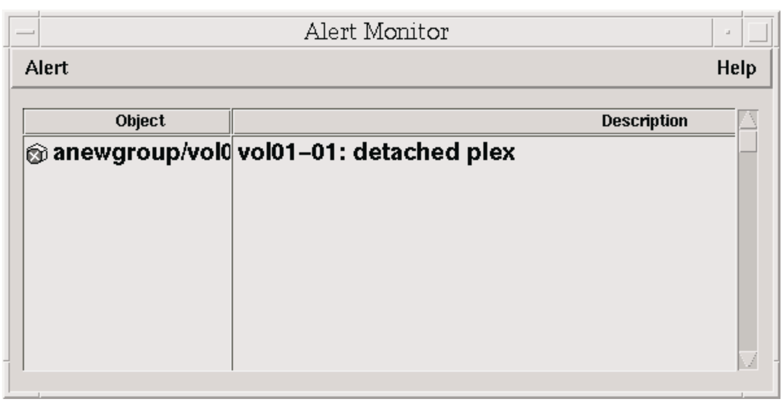

To display the Alert Monitor window either:

- Click on Alert in the toolbar
- Choose Alerts from the Window menu
- Click on the Alert icon on the Status bar

To view the properties of an object with an alert, select the object and choose Object Properties from the Alert menu. You can also access the object Properties dialog box by right-clicking and choosing Properties from the pop-up menu or by double-clicking on the object.

B.3.6 Object Table Copy Window

A copy of the Object Table allows you to view different parts of your system at the same time. The windows are dynamic, so updates to the system are reflected in all windows. To display the window copy, click on Table in the toolbar or, from the Window menu, choose Copy Object Table.

B.3.7 Object Search Window

The Object Search window searches the system for objects that match the specified search criteria. The Object Search window (Figure B–6) contains a set of tabbed pages that display search options for a particular type of object. You can select the type of objects to search for by clicking on the appropriate tab label. The search takes place only on objects of the type you have selected.

To display the Object Search window, click on Search in the toolbar or, from the Window menu, choose Search.

Figure B–6: Object Search Window

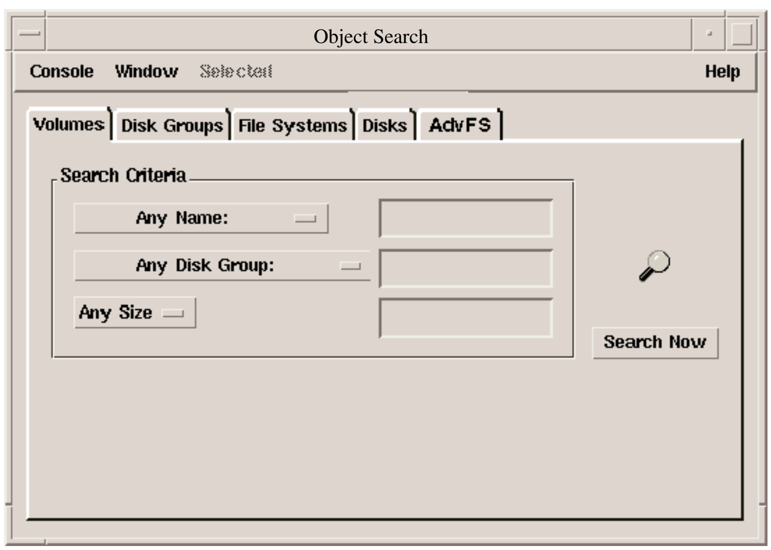

Specify the search criteria by selecting a qualifier from the drop-down menu and then entering a value to search for or exclude. For example, in the Volumes tab, you can select the Name Does Not Contain qualifier and enter the string “swap.” This excludes all volume names that contain the word “swap” from the search results. If you enter criteria in more than one field, the search results reflect only the items that match all criteria (Boolean AND).

The table in the bottom half of the Object Search window displays objects and their properties that match the search criteria. (If you do not see the search results, drag the bottom edge of the window to enlarge the display.) Objects displayed in the Object Search window are monitored and removed from the window if they no longer meet the current search criteria.

The Object Search window has menus similar to those in the Main window menu. The Window menu opens other windows or a copy of the current search-results table. The context-sensitive Selected menu accesses tasks or properties for an object selected in the table. You can right-click on an object to access a context-sensitive pop-up menu. To close the Object Search window, choose Close from the Console menu.

B.3.8 Task Request Monitor Window

The Task Request Monitor window (Figure B–7) shows LSM and other tasks that the Storage Administrator performed in the current session (and any other sessions running on the system).

To display the Task Request Monitor window, click on Task in the Storage Administrator Main Window, or from the Window menu, choose Tasks.

Figure B–7: Task Request Monitor Window

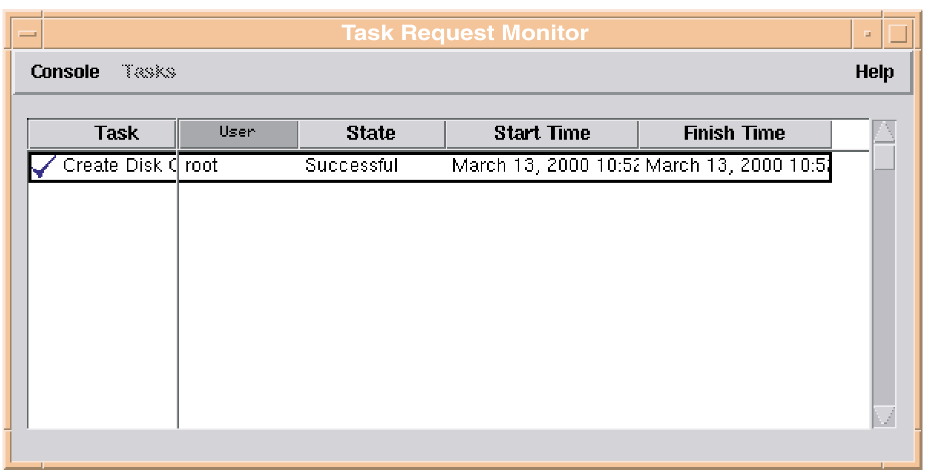

Each task is listed with properties such as the user who performed the task, the task status, and the start and finish times.

You can view the low-level commands used to perform a task and see any relevant error messages for failed tasks by displaying the Properties window. To display the Task Properties window, select a task and from the Tasks menu, select Properties. You can copy commands from the Executed Commands field of the Tasks Properties window to the command line or to a script file.

To remove finished tasks and to close the window, from the Console menu choose Remove Finished Tasks.

B.4 Storage Administrator Shortcuts

Shortcuts let you more efficiently perform operations with the GUI.

B.4.1 Sorting Objects

To sort the objects in a table column, click on the column heading. To reverse the order of the objects, click on the column heading again. The sort order cannot be saved with other user preferences.

You can sort entries in the Object Table, the Command Launcher, the Object Search window, and the Task Request Monitor window.

B.4.2 Clearing an Alert

To acknowledge and clear an Alert icon displayed on the Status bar, choose Clear Alert Status from the Options menu.

B.4.3 Keyboard Shortcuts

You can use the keyboard shortcuts in Table B–1 and Table B–2 instead of menu commands.

The shortcuts in Table B–1 work from anywhere in the Storage Administrator window.

Table B–1: Keyboard Shortcuts

Keystrokes	Action
Ctrl Shift V	Create a volume
Ctrl G	Create a disk group
Ctrl F	Create a file system
Ctrl Z	Resize the selected object
Ctrl N	Rename the selected object
Ctrl Shift R	Remove the selected object
Ctrl P	Show the properties of the selected object
Ctrl L	Show the layout of the selected volume (graphical view)

The shortcuts in Table B–2 work only in the Main window.

Table B–2: Main Window Keyboard Shortcuts

Keystrokes	Action
Ctrl R	Open the Task Request Monitor window
Ctrl A	Open the Alert Monitor window
Ctrl S	Open the Object Search window
Ctrl C	Close the window

B.4.4 Docking the Toolbar and Command Launcher

To separate the toolbar from the Main window, place the pointer over the toolbar handle (the thin bar next to the toolbar) and drag the toolbar outside the window. You can also use the toolbar handle to move the toolbar to the bottom, side, or top of the Main window.

You can separate or attach the Command Launcher and the Main window. By default, they are attached and the Command Launcher is not shown.

- To separate the Command Launcher from the Main window, choose Preferences from the Options menu. In the Preferences dialog box, choose the Main Window tab. Click on Show Command Launcher and then click on Dock Command Launcher.
- To attach the Command Launcher and the Main window, choose Preferences from the Options menu. In the Preferences dialog box, choose the Main Window tab. Click on Dock Command Launcher.

C

Using the voldiskadm Menu Interface

This appendix describes the `voldiskadm` menu interface that you can use to perform LSM disk and disk group operations. The menus are easy to use and provide information about each step to help you decide the correct response for each prompt.

C.1 Starting the voldiskadm Menu Interface

To start the `voldiskadm` menu interface, enter:

```
# voldiskadm
```

An interactive menu (Figure C–1) is displayed.

Figure C–1: Main Menu for the LSM voldiskadm Interface

```
Logical Storage Manager Support Operations
Menu: VolumeManager/Disk

 1      Add or initialize one or more disks
 2      Encapsulate one or more disks
 3      Remove a disk
 4      Remove a disk for replacement
 5      Replace a failed or removed disk
 6      Mirror volumes on a disk
 7      Move volumes from a disk
 8      Enable access to (import) a disk group
 9      Remove access to (deport) a disk group
 10     Enable (online) a disk device
 11     Disable (offline) a disk device
 12     Mark a disk as a spare for a disk group
 13     Turn off the spare flag on a disk
 14     Recover plexes and volumes after disk replacement
 list   List disk information

 ?      Display help about menu
 ??     Display help about the menuing system
 q      Exit from menus

Select an operation to perform:
```

C.2 Disk Management

This section describes the disk management tasks available with the `voldiskadm` menu interface.

C.2.1 Initializing a Disk

Disk initialization identifies a disk to LSM and prepares the disk for LSM use. This operation involves installing a disk header and writing an empty configuration database on the disk. A disk access record is created for the disk, unless such a record already exists.

Note

A disk must contain a disk label before you can initialize it for LSM use.

To initialize a disk for use with the LSM software:

1. At the main menu prompt, select menu item 1 to run the Add or initialize one or more disks operation.

   The Add or initialize disks screen appears.

2. Enter the disk access name of the disk to be added. If you do not know the access name of the disk you want to add, enter the letter `l` or enter the word `list` at the prompt. LSM displays a list of the known disks on the system. For example:

   ```
   Select disk devices to add:

   [<space-separated disk list>,<disk>,list,q,?] list

   DEVICE       DISK         GROUP        STATUS
   dsk2         dsk2         rootdg       online
   dsk3         dsk3         rootdg       online
   dsk4         dsk4         rootdg       online
   dsk5         dsk5         rootdg       online
   dsk6         -            -            online
   dsk7         -            -            online
   dsk8         -            -            unknown
   .
   .
   .
   Select disk devices to add:
   [<space-separated disk list>,<disk>,list,q,?]
   ```

3. Enter the name of the disk group you want the disk to be a part of:

   ```
   Which disk group [<group>,none,list,q,?] (default: rootdg)
   ```

   You can:

   - Press Return to accept the default disk group name, rootdg.

- Specify the name of a disk group to add the disk to an existing disk group.
- Create a new disk group and add the disk to it. To create a new disk group, enter a disk group name that does not yet exist.
- Specify no disk group and leave the disk available for use by future add or replacement operations. To leave the disk available for future use, enter a disk group name of none. Enter none if:
  – The disk group you want the disk to be a part of does not exist yet.
  – You want to keep this disk available as a spare to be used as a replacement disk.

4. Depending on your response to the Which disk group prompt, LSM displays one of the following screens.
   – If you entered none, LSM displays the following prompts:

```
Which disk group [<group>,none,list,q,?] (default: rootdg) none

The disk will be initialized and left free for use as a replacement
disk.

dsk8

Continue with operation? [y,n,q,?] (default: y) Return

The following disk device has a valid disk label, but does not appear to
have been initialized for the Logical Storage Manager.  If there is
data on the disk that should NOT be destroyed you should encapsulate
the existing disk partitions as volumes instead of adding the disk
as a new disk.

dsk8
Initialize this device? [y,n,q,?] (default: y) Return
Initializing device dsk8.
```

   – If you selected rootdg as the disk group, LSM displays the following prompts:

```
Which disk group [<group>,none,list,q,?] (default: rootdg) Return

The default disk name that will be assigned is:

disk01

Use this default disk name for the disk? [y,n,q,?] (default: y) Return

Add disk as a spare disk for rootdg? [y,n,q,?] (default: n) Return

The selected disks will be added to the disk group rootdg with the
default disk names.

dsk8

Continue with operation? [y,n,q,?] (default: y) Return

The following disk device has a valid disk label, but does not appear to
```

```
have been initialized for the Logical Storage Manager.  If there is
data on the disk that should NOT be destroyed you should encapsulate
the existing disk partitions as volumes instead of adding the disk
as a new disk.

dsk8

Initialize this device? [y,n,q,?] (default: y)  Return

Initializing device dsk8.

 Adding disk device dsk8 to disk group rootdg with disk
 name disk01.
```

5. Press Return to continue.
   - If LSM successfully completes the disk initialization, the following prompt appears:

     ```
     Add or initialize other disks? [y,n,q,?] (default: n)
     ```

   - If the `fstype` column in the disk's disk label information contains anything other than unused, LSM displays a message asking you to confirm or cancel the operation:
     - If the `fstype` in the disk label of the specified partition or an overlapping partition is set, LSM displays a warning message to inform you that initializing the disk might destroy existing data.
     - If you are sure that the disk partition has no valid data and that the partition can be added to LSM, you can ignore the warning message and answer `y` to the prompt. The `voldiskadm` utility then proceeds to initialize the disk partition and add it to LSM.
     - If the disk cannot be initialized because the specified partition or an overlapping partition on the disk is open (that is, a partition is actively in use by UFS, AdvFS, LSM, or swap), the initialization process fails and `voldiskadm` issues an error message informing you of the problem.
6. Press Return to return to the main menu.

C.2.2 Displaying Disk Information

To find information about disks available on the system:

1. At the main menu prompt, enter the letter `l` or enter the word `list` to display a list of disks available on the system.

   LSM displays a list of devices similar to the following and prompts you to enter the address (name) of the disk for which you want to obtain detailed information:

   ```
   Select an operation to perform: list

   List disk information
   Menu: VolumeManager/Disk/ListDisk
   ```

```
Use this menu operation to display a list of disks.  You can
also choose to list detailed information about the disk at
a specific disk device address.

Enter disk device or "all" [<address>,all,q,?] (default: all) Return

DEVICE       DISK         GROUP        STATUS
dsk2         dsk2         rootdg       online
dsk3         dsk3         rootdg       online
dsk4         dsk4         rootdg       online
dsk5         dsk5         rootdg       online
dsk6         -            -            online
dsk7         -            -            online
dsk8         disk01       rootdg       online
.
.
.
```

2. Enter the addresses (names) of all the disks for which you want detailed information (in this case, dsk8):

```
Device to list in detail [<address>,none,q,?] (default: none) dsk8

Device:    dsk8
devicetag: dsk8
type:      sliced
hostid:    servername
disk:      name=disk01 id=922907065.1771.servername
group:     name=rootdg id=921709207.1025.servername
flags:     online ready autoimport imported
pubpaths:  block=/dev/disk/dsk8g char=/dev/rdisk/dsk8g
privpaths: block=/dev/disk/dsk8h char=/dev/rdisk/dsk8h
version:   n.n
iosize:    min=512 (bytes) max=32768 (blocks)
public:    slice=6 offset=16 len=4106368
private:   slice=7 offset=0 len=4096
update:    time=922907069 seqno=0.5
headers:   0 248
configs:   count=1 len=2993
logs:      count=1 len=453
Defined regions:
config   priv     17-   247[   231]: copy=01 offset=000000 enabled
config   priv    249-  3010[  2762]: copy=01 offset=000231 enabled
log      priv   3011-  3463[   453]: copy=01 offset=000000 enabled

List another disk device? [y,n,q,?] (default: n)
```

3. Press Return to return to the main menu.

C.2.3 Adding a Disk to a Disk Group

You might want to add a new disk to an existing disk group. Perhaps the current disks have insufficient space for the project or work group requirements, especially if these requirements have changed.

To add a disk to a disk group:

1. Follow the instructions documented in steps 1 and 2 in Section C.2.1.

2. When the add disk operation adds a disk to a disk group, LSM checks whether the disk is already initialized. If the disk is initialized, LSM displays the following screen and prompts you to reinitialize the disk:

```
Which disk group [<group>,none,list,q,?] (default: rootdg) dg1

    The default disk name that will be assigned is:

    dg101

  Use this default disk name for the disk? [y,n,q,?] (default: y)  Return

  Add disk as a spare disk for dg1? [y,n,q,?] (default: n)  Return

    The selected disks will be added to the disk group dg1 with the
    default disk names.

    dsk8

    Continue with operation? [y,n,q,?] (default: y)  Return

      The following disk device appears to have been initialized already.
      The disk is currently available as a replacement disk.

      dsk8

    Use this device? [y,n,q,?] (default: y)  Return

    The following disk you selected for use appears to already have
    been initialized for the Logical Storage Manager.  If you are
    certain the disk has already been initialized for the Logical
    Storage Manager, then you do not need to reinitialize the disk
    device.

dsk8

  Reinitialize this device? [y,n,q,?] (default: y)
```

Use the information in the following table to determine whether or not you should reinitialize the disk.

If	Then
The disk is new	Initialize the disk before placing it under LSM control.
The disk was previously in use and contains useful data	Do not initialize the disk. Instead, use the LSM encapsulation function to place the disk under LSM control while still preserving the existing data.
The disk was previously in use but it does not contain useful data	Initialize the disk before placing the disk under LSM control.

C.2.4 Moving Volumes from a Disk

Before you disable or remove a disk, you might want to move the data from that disk to other disks on the system. Use this operation immediately prior to removing a disk, either permanently or for replacement (Section C.2.5).

Note

Simply moving volumes off a disk without also removing the disk does not prevent volumes from being moved onto the disk by future operations. For example, two consecutive move operations could move volumes from one disk to another and then back.

Make sure the other disks in the disk group have sufficient space available.

To move volumes from a disk:

1. Select item 7 from the main menu.
2. From the Move volumes from a disk screen, enter the name of the disk whose volumes you want to move:

```
Use this menu operation to move any volumes that are using a
disk onto other disks.  Use this menu immediately prior to
removing a disk, either permanently or for replacement.  You can
specify a list of disks to move volumes onto, or you can move the
volumes to any available disk space in the same disk group.

    NOTE:  Simply moving volumes off of a disk, without also removing
           the disk, does not prevent volumes from being moved onto
           the disk by future operations.  For example, using two
           consecutive move operations may move volumes from the
           second disk to the first.

  Enter disk name [<disk>,list,q,?] dsk5
```

3. Enter the name of the disk that the volumes should be moved to:

```
You can now specify a list of disks to move onto.  Specify a list
of disk media names (e.g., rootdg01) all on one line separated by
blanks.  If you do not enter any disk media names, then the volumes
will be moved to any available space in the disk group.

Enter disks [<disk ...>,list,q,?] dsk4

    Requested operation is to move all volumes from disk dsk5 in
    group rootdg.

      NOTE: This operation can take a long time to complete.

Continue with operation? [y,n,q,?] (default: y)  Return
```

   As LSM moves the volumes from the disk, it displays the status of the operation:

```
Move volume v1 ...
```

When the volumes have all been moved, LSM displays the following:

```
Evacuation of disk dsk5 is complete.
Move volumes from another disk? [y,n,q,?] (default: n)
```

4. Press Return to return to the main menu.

C.2.5 Removing a Disk from a Disk Group

This operation involves removing the LSM disk associated with the selected partitions from LSM control by removing the associated disk access records. The `voldiskadm` menu interface provides two methods for removing disks from a disk group:

- Menu item 3, Remove a disk—Removes a disk completely from LSM control and does not retain the disk name.
- Menu item 4, Remove a disk for replacement—Removes a failed disk and retains the disk name so it can be replaced with another disk.

See Section C.2.4, which describes how to move data from a disk to another disk on the system, and Section C.2.6, which describes how to replace a failed or removed disk.

Note

You must disable the disk group before you can remove the last disk in that group. Disabling a disk group, also referred to as deporting a disk group, is described in Section C.3.2.

C.2.5.1 Removing a Disk Without Replacement

To remove a disk from its disk group:

1. Select item 3 from the main menu.

   LSM displays the Remove a disk screen.

2. Enter the name of the disk to be removed (in this case, dsk5):

```
Use this operation to remove a disk from a disk group.  This
operation takes, as input, a disk name.  This is the same name
that you gave to the disk when you added the disk to the disk
group.

  Enter disk name [<disk>,list,q,?] dsk5
```

3. Confirm the operation by pressing Return at the verification prompt:

```
Requested operation is to remove disk dsk5 from group rootdg.

  Continue with operation? [y,n,q,?] (default: y) Return
```

   LSM removes the disk from the disk group and then displays the following when the operation has completed:

```
Removal of disk dsk5 is complete.

  Remove another disk? [y,n,q,?] (default: n)
```

4. Press Return to return to the main menu.

C.2.5.2 Removing a Disk for Replacement

You might occasionally need to replace a disk in a disk group. This operation involves initializing the disk for LSM use and replacing the old disk and associated disk media records with the new disk and its information.

To replace a disk while retaining the disk name:

1. Select item 4 from the main menu.

   LSM displays the Remove a disk for replacement screen:

   ```
   Use this menu operation to remove a physical disk from a disk
   group, while retaining the disk name.  This changes the state
   for the disk name to a "removed" disk.  If there are any
   initialized disks that are not part of a disk group, you will be
   given the option of using one of these disks as a replacement.


     Enter disk name [<disk>,list,q,?]
   ```

2. Enter the name of the disk to be replaced if you know it. Otherwise, enter the letter `l` for a list of disks, and then enter the disk name at the prompt.

   - If the operation does not need to disable a volume to replace the disk, the following is displayed:

     ```
     The following volumes will lose mirrors as a result of this
     operation:

     v1

      No data on these volumes will be lost.
     ```

   - If the operation must disable a volume to replace the disk, the following is displayed:

     ```
     The following volumes will be disabled as a result of this
      operation:

             tst

       These volumes will require restoration from backup.

     Are you sure you want to do this? [y,n,q,?] (default: n)
     ```

3. Confirm the operation by entering **y** and pressing Return.

4. If there are any initialized disks available that are not part of a disk group, LSM displays the following screen and gives you the option of using one of these disks as a replacement. Select the replacement disk from the list provided. Press Return if you want to use the default disk.

```
The following devices are available as replacements:

        dsk5

      You can choose one of these disks now, to replace dsk4.
      Select "none" if you do not wish to select a replacement disk.

    Choose a device, or select "none"
    [<device>,none,q,?] (default: dsk5) Return
```

5. Confirm the operation by pressing Return:

```
Requested operation is to remove disk dsk4 from group rootdg.
The removed disk will be replaced with disk device dsk5.

Continue with operation? [y,n,q,?] (default: y) Return
```

   When LSM successfully replaces the disk, LSM displays the following screen:

```
Removal of disk dsk4 completed successfully.

Proceeding to replace dsk4 with device dsk5.

Disk replacement completed successfully.

Remove another disk? [y,n,q,?] (default: n)
```

6. Press Return to return to the main menu.

C.2.6 Replacing a Failed or Removed Disk

To replace a disk that you removed with the Remove a disk for replacement menu operation or a disk that failed during use:

1. Select item 5 from the main menu.
2. Enter the name of an uninitialized disk to replace the failed or removed disk:

```
Replace a failed or removed disk
Menu: VolumeManager/Disk/ReplaceDisk

  Use this menu operation to specify a replacement disk for a disk
  that you removed with the "Remove a disk for replacement" menu
  operation, or that failed during use.  You will be prompted for
  a disk name to replace and a disk device to use as a replacement.
  You can choose an uninitialized disk, in which case the disk will
  be initialized, or you can choose a disk that you have already
  initialized using the Add or initialize a disk menu operation.

Select a removed or failed disk [<disk>,list,q,?] list

  Disk group: rootdg

  DM NAME        DEVICE       TYPE     PRIVLEN  PUBLEN   STATE

  dm dsk5        -            -        -        -        NODEVICE

  Select a removed or failed disk [<disk>,list,q,?] dsk5

  Select disk device to initialize [<address>,list,q,?] list
```

```
DEVICE       DISK         GROUP        STATUS
dsk2         dsk2         rootdg       online
dsk3         dsk3         rootdg       online
dsk4         dsk4         rootdg       online
dsk5         -            -            online
dsk10        -            -            unknown


Select disk device to initialize [<address>,list,q,?] dsk10

  The following disk device has a valid disk label, but does not appear to
  have been initialized for the Logical Storage Manager.  If there is
  data on the disk that should NOT be destroyed you should encapsulate
  the existing disk partitions as volumes instead of adding the disk
  as a new disk.

  dsk10

Initialize this device? [y,n,q,?] (default: y)  Return

  The requested operation is to initialize disk device dsk10 and
  to then use that device to replace the removed or failed disk
  dsk5 in disk group rootdg.

Continue with operation? [y,n,q,?] (default: y)  Return
```

LSM displays the following success screen:

```
Replacement of disk dsk5 in group rootdg with disk device
    dsk10 completed successfully.

 Replace another disk? [y,n,q,?] (default: n)
```

3. Press Return to return to the main menu.

C.2.7 Disabling a Disk

This operation places the disk access record in an offline state. During searches for disk IDs or members of a disk group, offline disks are ignored.

To disable a disk:

1. Select item 11 from the main menu.
2. Select the disk you want to disable (in this case, dsk3):

```
Use this menu operation to disable all access to a disk device
by the Logical Storage Manager.  This operation can be applied
only to disks that are not currently in a disk group.  Use this
operation if you intend to remove a disk from a system without
rebooting.

   NOTE:  Many systems do not support disks that can be removed from
          a system during normal operation.  On such systems, the
          offline operation is seldom useful.

 Select a disk device to disable [<address>,list,q,?] dsk3
```

LSM disables the disk and then prompts you to disable another device:

```
Disable another device? [y,n,q,?] (default: n)
```

3. Press Return to return to the main menu.

C.3 Disk Group Management

This section describes how to manage disk groups with the `voldiskadm` menu interface.

C.3.1 Importing a Disk Group

If you want to move a disk group from one system to another, you must first disable (deport) it on the original system (Section C.3.2), then move the disks between systems and enable (import) the disk group on the target system.

Note

If two hosts share a SCSI bus, make sure that the other host failed or deported the disk group. If two hosts import a disk group at the same time, the disk group will be corrupted and become unusable.

To enable access by the system or cluster to a disk group, or to move a disk group from one system to another:

1. Select item 8 from the main menu.
2. From the Enable access to (import) a disk group screen, select the name of the disk group to import:

```
Use this operation to enable access to a disk group.  This can be
used as the final part of moving a disk group from one system to
another.  The first part of moving a disk group is to use the
"Remove access to (deport) a disk group" operation on the
original host.

A disk group can be imported from another host that failed without
first deporting the disk group.  Be sure that all disks in the disk
group are moved between hosts.

If two hosts share a SCSI bus, be very careful to ensure that the
other host really has failed or has deported the disk group.  If
two active hosts import a disk group at the same time, the disk
group will be corrupted and will become unusable.

Select disk group to import [<group>,list,q,?] (default: list)  Return

        GROUP: dg1 (id: 921709259.1071.servername)
        DEVICES:
                 dsk7

        GROUP: dg1 (id: 922382892.1625.servername)
        DEVICES:
                 dsk9

        GROUP: dg1 (id: 922908695.1779.servername)
        DEVICES:
                 dsk6
```

```
                dsk8

  Select disk group to import [<group>,list,q,?] (default: list) dg1
```

When the import is complete, LSM displays the following success screen:

```
The import of dg1 was successful.

Select another disk group? [y,n,q,?] (default: n)
```

3. Press Return to return to the main menu.

C.3.2 Deporting a Disk Group

Deport a disk group if you want to:

- Move the disks in a disk group to another system
- Use all the disks remaining in a disk group for some new purpose

You cannot deport the rootdg disk group.

Note

For removable disk devices on some systems, it is important to disable all access to the disk before removing the disk.

To deport a disk group:

1. Select item 9 from the main menu.
2. From the Remove access to (deport) a disk group screen, enter the name of the disk group to deport:

```
Use this menu operation to remove access to a disk group that
is currently enabled (imported) by this system.  Deport a disk
group if you intend to move the disks in a disk group to another
system.  Also, deport a disk group if you want to use all of the
disks remaining in a disk group for some new purpose.

You will be prompted for the name of a disk group.  You will
also be asked if the disks should be disabled (offlined).  For
removable disk devices on some systems, it is important to
disable all access to the disk before removing the disk.

Enter name of disk group [<group>,list,q,?] (default: list) Return

GROUP        DISK/VOLUME         DEVICE/STATE LENGTH

dg1          disk   dg101        dsk8         4106368
dg1          disk   dsk6         dsk6         4109440

Enter name of disk group [<group>,list,q,?] (default: list) dg1

The requested operation is to disable access to the removable
disk group named dg1.  This disk group is stored on the
following disks:

          dsk6 on device dsk6
```

```
          dg101 on device dsk8

You can choose to disable access to (also known as "offline")
these disks.  This may be necessary to prevent errors if
you actually remove any of the disks from the system.

Disable (offline) the indicated disks? [y,n,q,?] (default: n)  Return

Continue with operation? [y,n,q,?] (default: y)
```

3. Confirm the operation by pressing Return.

   When the disk group is deported, the following is displayed:

```
Removal of disk group dg1 was successful.

Disable another disk group? [y,n,q,?] (default: n)
```

4. Press Return to return to the main menu.

C.4 Mirror Volume Management

You can use the voldiskadm menu interface to add a mirror to a volume with no mirrors. You cannot use the voldiskadm menu interface to add mirrors to volumes that already have mirrors or that are comprised of more than one subdisk.

To mirror volumes on a disk:

1. Make sure that the target disk has an amount of space equal to or greater than the originating disk.
2. Select item 6 from the main menu.
3. In the Mirror volumes on a disk screen, enter the name of the disk whose volumes you want to mirror:

```
This operation can be used to mirror volumes on a disk.  The volumes
can be mirrored onto another disk or onto any available disk space.
Volumes will not be mirrored if they are already mirrored or contain
more than one subdisk.

Mirroring the root and swap volumes from the boot disk will produce a
disk that can be used as an alternate boot disk.

At the prompt below, supply the disk media name containing the
volumes to be mirrored.

Enter disk name [<disk>,list,q,?] dsk5
```

4. Enter the target disk name. Volumes can be mirrored onto another disk or onto any available disk space.

```
You can choose to mirror volumes from disk dsk5 onto any
available disk space, or you can choose to mirror onto a specific
disk.  To mirror to a specific disk, select the name of that disk.
To mirror to any available disk space, select "any".

Enter destination disk [<disk>,list,q,?] (default: any) dsk4
```

LSM displays a verification screen.

5. Press Return to make the mirror.

```
The requested operation is to mirror all volumes on disk dsk5
in disk group rootdg onto available disk space on disk dsk4.

There is space already allocated on disk dsk4.  If you don't
want to mirror onto this disk, enter "n" at the next prompt and
restart this operation from the beginning.

NOTE: This operation can take a long time to complete.

Continue with operation? [y,n,q,?] (default: y)  Return
```

LSM displays the status of the operation as it performs the mirroring:

```
 Mirroring of disk dsk5 is complete.
```

When the mirroring operation is complete, LSM prompts you to mirror volumes on another disk.

```
Mirror volumes on another disk? [y,n,q,?] (default: n)
```

6. Press Return to return to the main menu.

C.5 Exiting the voldiskadm Menu Interface

When you have completed all your disk administration activities, exit the `voldiskadm` menu interface by selecting menu option `q` from the main menu.

D

The Visual Administrator Interface

The Visual Administrator, also called `dxlsm`, is the graphical user interface for the LSM software. This interface is designed primarily for disk and volume operations, but also provides a limited set of file system operations.

This appendix provides an overview of Visual Administrator features and use.

D.1 Starting the Visual Administrator

To start the LSM Visual Administrator, you must be logged into an account that has root privileges.

To start `dxlsm`, enter:

```
# dxlsm
```

The system displays the following message in a pop up window:

```
dxlsm is coming up, please wait.
```

When `dxlsm` starts, two LSM windows display:

- The main LSM Visual Administrator window
- The View of rootdg window

D.1.1 The Visual Administrator Main Window

The Visual Administrator main window contains a Menu Bar and a set of buttons as shown in Figure D–1. If you have RAID subsystems installed, the main window displays an additional Subsystems button.

Figure D–1: Visual Administrator Main Window

To display and manipulate different parts of the physical and logical storage systems, click on View. Each view window title includes the name of the machine on which the session is running. The main window contains a button for every view on the system.

The Visual Administrator has two types of views: default and user-created. You cannot remove or rename default views.

D.1.1.1 Default Views

From the main window, click on the following view buttons to access the default view window:

Table D–1: Accessing the Default View Window

View Buttons	Window	Access
Disks	View of disks	Displays all physical disks on the system
Volumes	View of volumes	Displays all volumes, plexes, and associated subdisks on the system
World	View of world	Displays everything on the system including physical and LSM disks, volumes, and other objects
rootdg	View of rootdg	Displays everything in the default disk group, rootdg, including LSM disks, volumes, and other objects

D.1.1.2 User-Created Views

A user-created view is a view window defined to focus on a part of a system. For example, you can create a view window for each disk group. Create a new view window with the Views menu on the main window. Creating this window will place a new button on the window. Once the view is created, you can add icons by selecting an icon from another view and either using the Icon menu Create Icons option or dragging and dropping the selected icon.

User-created views thus contain copies of icons from default views. Operations performed on these icon copies are reflected in the default views that display the affected icons. However, icons that appear in user-created views are not always updated when the icons are altered in the default view.

D.1.2 The View of rootdg Window

Immediately after the main window appears, the View of rootdg window appears. This view displays icons representing everything that is in the rootdg disk group. Whenever possible, perform operations in the View of rootdg window or in another disk group view.

Figure D–2 shows a View of rootdg window that contains two volumes.

Figure D–2: View of rootdg Window

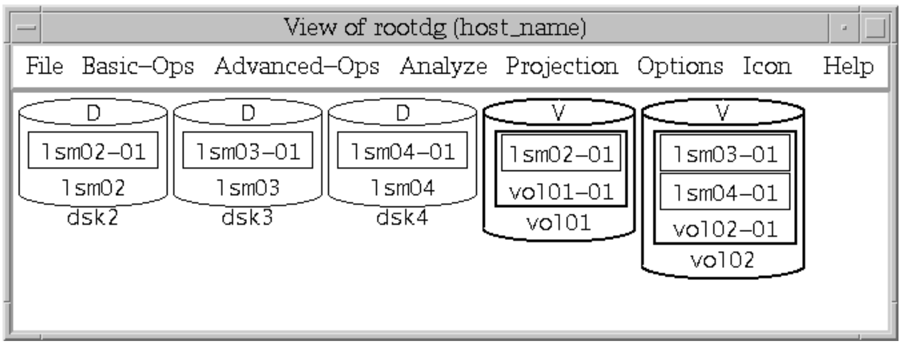

D.2 Mouse Buttons

A two- or three-button mouse is required to use the Visual Administrator. Table D–2 describes the mouse buttons. Unless otherwise stated, all directions to select or click on an item refer to the left mouse button (MB1). Right-click refers to the right mouse button for both two- and three-button mice.

Table D–2: Default Mouse Buttons

Virtual Mouse Button	3-Button Access	2-Button Access	Function
MB1	Left	Left	Selects a single icon.
MB2	Middle	Ctrl–Left	Selects either one or multiple icons simultaneously.
MB3	Right	Right	Displays either the properties form or the analysis statistics form for that object, depending on whether the icon is undergoing analysis.
Shift–MB1	Shift–Left	Shift–Left	Toggles between minimizing and maximizing an icon.
Shift–MB2	Shift–Middle	Ctrl–Right	Toggles between starting and stopping projection on the selected icon.
Shift–MB3	Shift–Right	Shift– Right	Displays the properties form for the object, regardless of whether analysis is in effect.

Deselect an icon by positioning the pointer over that icon and clicking MB2. This works regardless of which mouse button you use to select the icon.

Note

The examples in this document assume that you are using a three-button mouse, set up according to Table D–2. It is possible to redefine mouse buttons (using the `xmodmap` command, for example). See your X Window System documentation for details.

D.3 Icons for LSM Objects

The Visual Administrator interface uses icons to represent LSM objects (volumes, plexes, subdisks, and disks.)

Disk groups are represented as view windows rather than icons. The icons representing LSM disks, volumes, and other objects belonging to a particular disk group are all displayed within the view of the disk group.

Note

Icons representing all elements of the system are displayed in the View of World window accessed by clicking on World on the Visual Administrator window.

D.3.1 Icon Characteristics

Figure D–3 shows how the Visual Administrator represents LSM objects.

Figure D–3: Icons That Represent LSM Objects

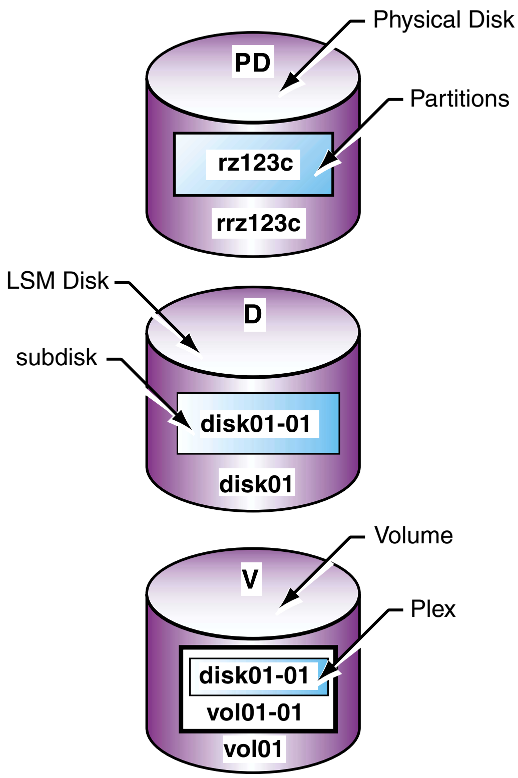

With some operations, icons are updated almost instantly to reflect the results of the operation just performed. During other operations, it might take time for a particular icon to update itself. While being updated, icons are prevented from accepting input or undergoing configuration changes. Inaccessible icons are grayed out.

D.3.2 Manipulating Icons

There are two ways to manipulate icons:

- Select–Operate. With this option you select an icon representing an LSM object (click on the icon or click MB2 on multiple icons) and perform the desired operation on that object by selecting from window menus.
- Drag and Drop. With this option you drag the icon of the selected object (hold down MB1 and move the mouse until the outline of the icon reaches the desired location) and drop it (release the mouse) elsewhere, such as on another object in a view window.

  When dropping an icon onto another icon, the dragged icon must be positioned so that the pointer (in the image of a hand) is directly over an unobscured portion of the icon on which it is to be dropped. Table D–3 describes drag and drop operations:

Table D–3: Drag and Drop Operations

Icon Type	Drop Location	Action
Free subdisk	View window	Creates a plex and associates the subdisk with the plex.
Free subdisk	LSM disk	Creates an identical subdisk on the disk.
Free subdisk	Plex	Associates the subdisk with the plex.
Associated subdisk	Free subdisk	Swaps the associated subdisk with the free subdisk. The free subdisk becomes associated and replaces the original subdisk, which is removed.
Associated subdisk	LSM disk	Creates an identical free subdisk on the LSM disk, then swaps the associated subdisk with the new free subdisk. The free subdisk becomes associated and replaces the original subdisk, which is removed.
Associated subdisk	View window	Dissociates the subdisk.
Associated plex	View window	Dissociates the plex.
Dissociated plex	User's view window	Copies the plex icon to the user's view.
Plex	Volume	Associates the plex to the volume.
Disk	User's view window	Copies the physical or LSM disk icon to the user's view.
Volume	User's view window	Copies the volume icon to the user's view.
LSM disk	Disk group view	Adds an LSM disk (corresponding to the slice, simple, or nopriv disk) to that disk group.

D.4 Pull-Down Menus

The Visual Administrator provides pull-down menus that provide access to various Visual Administrator features.

Menus are located in the Menu Bar just below the window's title.

D.5 Forms

The Visual Administrator uses forms to present textual information. These forms also provide useful information about existing objects and configurations.

There are two types of forms:

- General forms usually appear during menu-selected operations or setup requests and accept or require user input.
- Properties forms display detailed information about a specific object's characteristics, some of which can be modified directly. Access properties forms by clicking MB3 on the chosen icon. (If the icon is undergoing analysis, use Shift–MB3).

D.5.1 Fields

Many forms require information in order to proceed with an operation. If a required field in the form is either left blank or is incorrect, an error will result. Other fields already contain information (such as default values), which you can either alter or accept. Yet other fields are read-only; these fields beep if you attempt to change them.

D.5.2 Form Error Messages

Error messages are displayed if you select Apply with incorrect fields on a form. A message is printed at the bottom of the form just above the buttons, and you can correct the values for those fields. If the error cannot be corrected or the operation is no longer desired, select Cancel.

D.6 Error and Warning Messages

The Visual Administrator uses dialog boxes to present error or warning messages. When a message is displayed, you cannot proceed until you select one of the buttons displayed in the error dialog box. Some warning boxes announce that a prerequisite is not met and require you to acknowledge this by clicking Continue before retrying the operation.

D.7 Help Windows

You can access online help text from the Menu Bar of the main window and from the view windows. Help text is also available through the Help option in submenus or by clicking Help at the bottom of forms. The Help window contains information about the current window, menu, form, or operation.

At the bottom of each Help window is a SEE ALSO area that lists related Help topics. To access any of the listed Help topics, click on the appropriate words in the SEE ALSO list. The Help facility keeps track of the order in which Help topics are visited, so you can move between topics by selecting Previous or Next from the Menu Bar.

The Help menu in the Menu Bar of the Help window itself provides access to the following information:

Table D–4: Help Menu Options

General Help	Accesses Help text that includes general information on the Visual Administrator Help facility and how it is used.
Help Index	Access a complete listing of the available Help topics arranged in logical groupings. Once you select a topic from this list, you can access that topic directly from the SEE ALSO section of this Help window, which lists all topics alphabetically.

To close the Help window, select the Close option from the File menu. The record of help topics visited is retained.

E

Using the Visual Administrator Interface

This appendix shows you how to complete common LSM management tasks using the Visual Administrator.

E.1 Managing Volumes

The following sections provide information on menus and forms relating to volume management.

E.1.1 Volume Menus

Both the Basic-Ops and Advanced-Ops menus provide access to volume-related menus. Most menus provide a Help selection, which contains information about the items and operations listed in that particular menu.

E.1.1.1 Basic-Ops Menu

You access the Basic-Ops menu by selecting:

Basic-Ops → Volume Operations

This menu provides access to volume operations involving general volume maintenance. These operations use the automated approach to volume management.

The Volume Operations menu provides the following selections:

- Create
- Remove Volumes Recursively
- Add Plex
- Remove Plex
- Resize
- Snapshot
- Help

The following list describes these menu selections:

- Create

  Basic-Ops → Volume Operations → Create

This operation creates a simple or striped volume on one or more disks. You can select one or more disks on which to create the volume (providing that there is sufficient space on the disks). If no disks are specified, the LSM software automatically determines the disks to use based on available free space.

From the Create menu, select the type of volume to be created from a submenu listing two of the basic types of volumes:

Table E–1: Create Volume Menu

Type	Description
Simple	Creates a simple, concatenated volume whose subdisks are arranged sequentially and contiguously within a plex.
Striped	Creates a volume with data spread fairly evenly across multiple disks by way of striping. Stripes are relatively small, equally-sized fragments that are allocated alternately to the subdisks within a plex.
RAID 5	Creates a volume with data and parity spread evenly and alternately across subdisks within a plex.

To create a mirrored volume, create a simple or striped volume, then mirror it using the Add Mirror option.

Requirements:

– Only disks in the same disk group can be selected.

– Only LSM disks (disks under LSM control and assigned to a disk group) can be selected.

– If striping is to be in effect, at least two disks are required in order for the operation to succeed.

Forms: Simple Volume/FS Create Form and Striped Volume/FS Create Form.

- Remove Volumes Recursively

Basic-Ops → Volume Operations → Remove Volumes Recursively

This operation removes the selected volumes and deallocates all the disk space set aside for that volume. It automatically removes all underlying plexes and subdisks associated with the volume.

Note

This is a permanent operation and cannot be undone. If completed, it will be difficult or impossible to retrieve the data associated with that volume. For this reason, a confirmation window is presented if the selected volume is not ready for removal (that is, started or enabled).

Requirements:

– At least one volume icon must be selected.

– The selected volumes cannot contain a mounted file system.

- Add Plex

  Basic-Ops → Volume Operations → Add Mirror

  This operation adds a plex to the selected volume by associating a plex of the correct length to the volume. The plex effectively duplicates the information contained in the volume. Although a volume can have a single plex, at least two are required for mirroring.

  From the Add Mirror menu, select the type of plex to be added from a submenu listing two of the basic types of plexes:

Table E–2: Add Mirror Menu Options

Type	Description
Simple	Adds a simple, concatenated plex whose subdisks are arranged sequentially and contiguously.
Striped	Adds a plex whose data is striped (allocated evenly and alternately across each of its subdisks).

You can select disks for this operation. However, the number of selected disks must be sufficient to accommodate the layout type of both the existing volume and the plex to be added. If no disks are selected, the free space for the plex is allocated by the LSM software.

Requirements:

– A volume icon must be selected.

– For a striped plex, at least two disks other than those already in use by the volume must be available.

- Remove Plex

  Basic-Ops → Volume Operations → Remove Mirror

  This operation removes the selected plex, along with any associated subdisks.

  Requirements:

  – A plex icon must be selected.

  – The last valid plex in a started or enabled volume cannot be removed.

- Resize

Basic-Ops → Volume Operations → Resize

This operation resizes the selected volume. The volume can be increased to, increased by, reduced to, or reduced by a given length. This involves adding or removing disk space to or from the plexes associated with the volume.

If new disk space is needed during the resize, it is allocated as necessary; if space becomes unused, it is added to the free space pool.

Requirements:

– A volume icon must be selected.

– A volume containing a mounted file system cannot be reduced.

Form: Volume Resize Form

- Snapshot

Basic-Ops → Volume Operations → Snapshot

This operation backs up a volume by creating a snapshot image of that volume. This is a convenient way of performing backup with minimal interruption.

This operation creates a new volume that is a snapshot of an existing volume. This is done by creating a plex of the existing volume (creating and associating a plex) using disk space from the pool of free disk space. The plex is synchronized to the rest of the volume (this might take some time) and a separate (snapshot) volume is then created for that plex. The snapshot volume represents a consistent copy of the original volume at the time the snapshot was begun. The snapshot volume can be used to make a backup of the original volume without stopping it. After the backup is made, you can remove the snapshot volume without losing any data.

Note

For UFS volumes, it is recommended that you unmount the file system briefly to ensure the snapshot data on disk is consistent and complete.

From the Snapshot menu, a submenu allows you to first create the snapshot plex and then the snapshot volume:

Option	Description
Snapstart	Start the snapshot procedure by creating a snapshot plex within the volume to be backed up. It takes a variable amount of time to update the new plex, during which time the snapshot plex icon is grayed out.
Snapshot	At a convenient time (preferably after warning users to suspend activity briefly), create another volume for the snapshot plex. This portion of the procedure should take only seconds to complete.

Requirements:

– A volume icon must be selected.

– There must be sufficient free disk space to accommodate the snapshot volume.

Form: Snapshot Form

E.1.1.2 Advanced-Ops Menu

You access the Advanced-Ops menu selections by selecting:

Advanced-Ops → Volume

This menu provides access to assorted volume operations. These volume operations use the manual approach to volume management. The Volume menu provides the following selections:

- Create
- Remove Volumes
- Initialize Volumes
- Start Volumes
- Stop Volumes
- Resynchronize Volumes
- Set to Maint State
- Recover Volumes
- Help

The following list describes these menu selections:

- Create

  Advanced-Ops → Volume → Create

  This operation creates a volume. You can select one or more plexes to be associated with the new volume after creation.

  Form: Volume Create Form

- Remove Volumes

  Advanced-Ops → Volume → Remove Volumes

  This operation removes the selected volumes. If the selected volume is started, it must be stopped before it can be removed.

  Note

  This is a permanent operation and cannot be undone. Any plexes associated with the volume will be dissociated and left behind.

  Requirements:

  – At least one volume icon must be selected.

  – The volume must be stopped before it can be removed.

- Initialize Volumes

  Advanced-Ops → Volume → Initialize Volumes

  This operation initializes the selected volumes.

  From the Initialize volumes menu, you select the type of initialization from a submenu listing the following choices:

Option	Description
Active	This enables the selected volume and its associated plexes, and sets the state of all associated plexes to ACTIVE.
Enable	This enables the selected volume and its associated plexes, but leave the plex states as EMPTY.
Clean	This sets the state for all associated plexes of the selected volume to CLEAN. This can be applied only under limited circumstances.
Zero	This enables the selected volume and its associated plexes, then write zeroes over the entire volume. After the operation completes, all associated plexes are set to ACTIVE, assuming that there are no I/O errors.

  Requirements:

  – At least one volume icon must be selected.

  – The selected volume cannot have been previously initialized.

  – The selected volume should have at least one associated plex that is complete (or contiguous).

- Start Volumes

  Advanced-Ops → Volume → Start Volumes

This operation starts the selected volumes. A volume must be started before it can be accessed.

From the Start volumes menu, a submenu allows you to indicate whether all volumes or just those selected should be started:

Option	Description
Start	Start the selected volume, which must be startable.
Start All	Start all volumes in this disk group that can be started.

Requirements:

- At least one volume icon must be selected for the Start operation. No volume icons need to be selected for the Start All operation.
- A volume should be initialized before it can be started.

• Stop Volumes

Advanced-Ops → Volume → Stop Volumes

This operation stops the selected volumes. A volume that is stopped is inaccessible.

From the Stop volumes menu, a submenu allows you to indicate whether all volumes or just those selected should be stopped:

Option	Description
Stop	Stop the selected volume.
Stop All	Stop all volumes in this disk group.

Requirements:

- At least one volume icon must be selected for the Stop operation. No volume icons need to be selected for the Stop All operation.
- A volume must be started before it can be stopped.
- A volume that is in use or contains a mounted file system cannot be stopped.

• Resynchronize Volumes

Advanced-Ops → Volume → Resynchronize Volumes

This operation brings all plexes within the selected volumes up to date. Any plexes that are inconsistent are resynchronized to contain consistent data.

This operation might take some time depending on how large the plexes are and whether or not logging is enabled.

Requirements:

- At least one volume icon must be selected.
- The selected volumes must be started.

- Set to Maintenance State

  Advanced-Ops → Volume → Set to Maint State

  This operation sets the state of the selected volumes to a maintenance state. Refer to the `volume(8)` reference page for information on the maintenance state.

  Requirement: At least one volume icon must be selected.

- Recover Volumes

  Advanced-Ops → Volume → Recover Volumes

  This operation recovers the selected volumes. At least one volume icon must be selected.

E.1.2 Volume Forms

Some volume operations result in the appearance of forms, which must be completed in order for that operation to proceed. Most forms provide a Help button, which contains information relevant to the fields and other aspects of that particular form.

E.1.2.1 Basic-Ops Forms

The following forms are accessed via volume-related selections from the Basic-Ops menu:

- Simple Volume/FS Create Form

  Basic-Ops → Volume Operations → Create → Simple

  This form creates a concatenated volume and optionally creates a file system on the new volume. The form is divided into two sections, one for volume creation and the other for file system creation. Most of the form fields are already set to the defaults for the creation of a new volume; the file system fields are grayed out because the default is not to add a file system to the volume. The following tables describes the fields for this form.

  Most fields in this form are required; those that are optional are listed here. All fields in this form are read/write fields.

Field	Description
Volume name:	The name of the volume to be created. The name must be unique within this disk group. The maximum length of this field is 31 characters.
Volume size:	The desired volume size. The size should be entered as a number followed immediately by the letter `k`, `m`, or *s* to indicate kilobytes, megabytes, or sectors, respectively. If no unit is specified, the default is sectors. The volume size should be less than or equal to the available free space of the disks.
Usage Type:	The desired usage type. The `fsgen` type is the file system generic usage type, which assumes that the volume is being used by a file system. The `gen` type is the generic usage type, which makes no assumptions regarding the data content of the volume. The default is `fsgen`.
Create file system:	Indicates whether a file system is to be created. When you invoke this form from the Volume Operations menu, the default is not to create a file system (No). All fields below this field are accessible only when Yes is specified here.

The following fields apply only if the Create file system: field is set to Yes. Otherwise, these fields are inaccessible.

Field	Description
FS type:	UFS is the only currently supported files system type.
Mount file system:	Indicates whether the file system should be mounted after creation. If the answer is Yes (the default), a mount point must also be specified in the next field. All fields below this field are accessible only when Yes is specified here.
Mount point:	The desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator automatically creates it. This field is required if the file system is to be mounted.
Mount automatically:	Indicates whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/fstab`). The default is Yes.

- Striped Volume/FS Create Form

  Basic-Ops → Volume Operations → Create → Striped

  This form creates a concatenated volume and optionally creates a file system on the new volume. The form is divided into two sections, one for volume creation and the other for file system creation. Most of the form fields are already set to the defaults for the creation of a new volume; the file system fields are grayed out because the default is not to add a file system to the volume. The following table describes the fields for this form.

Most fields in this form are required; those that are optional are listed here. All fields in this form are read/write fields.

Field	Description
Volume name:	The name of the volume to be created. The name must be unique within this disk group. The maximum length of this field is 31 characters.
Volume size:	The desired volume size. The size should be entered as a number followed immediately by the letter `k`, `m`, or *s* to indicate kilobytes, megabytes, or sectors, respectively. If no unit is specified, the default is sectors. If the size is not wholly divisible by the stripe width, LSM will adjust the volume size up to the next even multiple in order to create the volume. For a striped volume, the volume size should be calculated as follows: *vol_size* = *stripe_width* * *number_of_stripes* * *n*, where *n* is a number greater than zero. The volume size should be less than or equal to the available free space of the disks.
Usage Type:	The desired usage type. The `fsgen` type is the file system generic usage type, which assumes that the volume is being used by a file system. The `gen` type is the generic usage type, which makes no assumptions regarding the data content of the volume. The default is `fsgen`.
Number of Stripes:	The number of stripes that the volume's plex is to have. This is effectively the number of disks on which the volume is to be created. If some number of disks have already been selected, that number of stripes appears in this field. This number corresponds to the number of disks across which data will be striped. If no number is specified, LSM selects an appropriate number (usually 2).
Stripe width:	The width of the stripes on the plex that this volume will have. The value specified might be optimized for the particular application. However, the default value for this field of 128 sectors is a good stripe width for most systems.
Create file system:	Indicates whether a file system is to be created. When you invoke this form from the Volume Operations menu, the default is not to create a file system (`No`). All fields below this field are accessible only when Yes is specified here.

The following fields apply only if you set the Create file system: field to Yes. Otherwise, these fields are inaccessible.

Field	Description
FS type:	UFS is the only currently supported file system type.
Mount file system:	Indicates whether the file system should be mounted after creation. If the answer is Yes (the default), you must also specify a mount point in the next field. All fields below this field are accessible only when Yes is specified here.

Field	Description
Mount point:	The desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator automatically creates it. This field is required if the file system is to be mounted.
Mount automatically:	Indicates whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/fstab`). The default is Yes.

- Volume Resize Form

Warning

File systems and other applications cannot currently resize their data when LSM resizes a volume, therefore shrinking a volume that contains data destroys the data. Therefore, use this operation only when a volume contains no valuable data.

Basic-Ops → Volume Operations → Resize

This form either grows or shrinks a volume using the Logical Storage Manager free space management resources. If new disk space is needed, it will be allocated as necessary; if space becomes unused, it will be added to the free space pool. The following table describes the fields in this form.

Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless they are listed as read-only.

Field	Description
Selected Volume:	This field displays the name of the volume to be resized. This field is read-only and cannot be changed.
Current size:	This field displays the current size of the volume to be resized. This field is read-only and cannot be changed.
Option:	The type of resize operation to be performed. This will determine whether the volume is grown or shrunk to a certain size, or grown or shrunk by a given amount. The default is Grow To.
Size/Amount:	Enter either the length to which or the amount by which the volume is to be resized. If Grow To or Shrink To is selected, this field should reflect the final size. If Grow By or Shrink By is selected, this field should reflect the amount by which the size should change. The new volume size should be less than or equal to the available free space of the disks.

- Snapshot Form

Basic-Ops → Volume Operations → Snapshot

This form creates a snapshot of the selected volume for backup purposes. The following table describes the fields for this form. Fields in this form are required.

Fields in this form are read/write fields, unless they are listed as read-only.

Field	Description
Selected Volume:	The name of the volume to be used as the snapshot source. This field is read-only and cannot be changed.
Snapshot name:	The name of the snapshot volume to be created as a backup. Although a default name appears in this field, a name that more closely resembles that of the selected volume should be used for easier association. The maximum length is 31 characters. The snapshot name must be unique.

Requirement: There must be sufficient free space to accommodate the snapshot volume.

E.1.2.2 Advanced-Ops Forms

The following forms are accessed via volume-related selections from the Advanced-Ops menu:

- Volume Create Form

  Advanced-Ops → Volume → Create

  This form creates a volume according to your specifications. The following table describes the fields for this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Volume name:	The name of the volume to be created. The name must be unique within this disk group. The maximum length of this field is 31 characters. The name specified for the volume must be unique within this disk group.
Usage Type:	The desired usage type. The `fsgen` type is the file system generic usage type, which assumes that the volume is being used by a file system. The `gen` type is the generic usage type, which makes no assumptions regarding the data content of the volume. The default is `fsgen`. This field is optional.
User:	The name of the user who will be the owner of this volume. This must be a valid user name on the system. The maximum length of this field is 64 characters.

Field	Description
Group:	The name of the group that will own this volume. This must be a valid group name on the system. The maximum length of this field is 64 characters.
Mode:	The permissions mode for the new volume. Only numbers of the correct format are valid in this field. The maximum length of this field is four characters.
Length:	The length of the volume. If no unit is specified, the default is sectors. Only positive numbers greater than zero are valid. This field is optional.
Plexes:	This field displays the number of plexes associated with the volume. If no plexes were selected prior to invoking this form, this field displays 0. This field is read-only and cannot be changed.
Read Policy:	The read policy that the volume adopts when deciding which plex to write to. These policies are distinguished as follows: – Round Robin—All plexes are read equally, in turn. – Preferred Plex—A particular plex is specified as the plex to be read whenever possible. The preferred plex will not be read in situations such as when that plex is detached due to I/O failure. – Based on plex layouts—All plexes are read equally and in turn, unless a striped plex is present, in which case the striped plex becomes the preferred plex. This option is the default and it typically gives the best read performance.
Preferred Plex:	The name of the preferred plex if the Preferred Plex read policy has been specified. The string in this field must be the name of a valid plex that is associated with this volume. This field is required if Preferred Plex is specified in the Read Policy: field.
Comment:	An appropriate comment for this volume. The maximum length of the comment is 40 characters. This field is optional.
Startup:	This field might contain an arbitrary string that is reserved for the user by usage-type utilities. The intention is that this field be used to store options that apply to the volume, such as for the start volumes operation. This is normally a comma-separated list of flag names and *option=value* pairs. This field is optional.
Logging:	Indicates whether logging is defined and supported on this volume. An undefined log type is included to support old versions of the Logical Storage Manager. The default is Don't Log.
Writeback:	Indicates whether the volume is to write back on read failure. If set to Yes, an attempt will be made to fix a read error from a participating plex. The default is Yes.

Field	Description
Putil0:	Permanent utility field 0. This is reserved for Logical Storage Manager use, but can be changed. The maximum length of all Putil fields is 14 characters. This field is optional.
Putil1:	Permanent utility field 1. This field is reserved, but can be changed. This field is optional.
Putil2:	Permanent utility field 2. This field is reserved, but can be changed. This field is optional.

E.1.2.3 Volume Properties Form

The following is the properties form that reveals the properties of a particular volume:

- Volume Properties Form

  You can access this form by clicking the MB3 on the desired volume icon. (If the volume icon is undergoing analysis, press Shift–MB3 instead.)

  This form provides detailed information on the attributes of a particular volume. The following table describes the fields in this form.

  The fields in this form are read/write fields, unless listed as read-only. Properties of the volume can be changed via this form by altering the current values in the appropriate read/write fields and then clicking on Apply.

Field	Description
Volume name:	The name of the volume. This name must be unique within this disk group. The maximum length of this field is 31 characters. This volume name can be changed by entering another name in this field.
Usage Type:	The volume usage type. The `fsgen` type is the file system generic usage type, which assumes that the volume is being used by a file system. The `gen` type is the generic usage type, which makes no assumptions regarding the data content of the volume.
Utility State:	The state that the volume is currently in. This should be either Started, Startable, or Unstartable. This field is read-only and cannot be changed.
User:	The name of the user who owns this volume. This must be a valid user name. The maximum length of this field is 64 characters.
Group:	The name of the group that will own this volume. This must be a valid group name. The maximum length of this field is 64 characters.

Field	Description
Mode:	The permissions mode for the volume. Only numbers of the correct format are valid in this field. The maximum length of this field is four characters.
Length:	The length of the volume. If no unit is specified, the default is sectors. Only positive numbers greater than zero are valid.
Plexes:	This field displays the number of plexes associated with the volume. If no plexes were selected prior to invoking this form, this field displays 0. This field is read-only and cannot be changed.
Read Policy:	The read policy that the volume adopts when deciding which plex to write to. These policies are distinguished as follows: – Round Robin—All plexes are read equally, in turn. – Preferred Plex—A particular plex is specified as the plex to be read whenever possible. The preferred plex will not be read in situations such as when that plex is detached due to I/O failure. – Based on plex layouts—All plexes are read equally and in turn, unless a striped plex is present, in which case the striped plex becomes the preferred plex. This option is the default and it typically gives the best read performance.
Preferred Plex:	The name of the preferred plex if the Preferred Plex read policy has been specified. The string in this field must be the name of a valid plex that is associated with this volume. This field applies only if Preferred Plex is specified in the Read Policy: field.
Comment:	A comment relevant to this volume. The maximum length of the comment is 40 characters.
Startup:	This field might contain an arbitrary string that is reserved for the user by usage-type utilities. The intention is that this field be used to store options that apply to the volume, such as for the start volumes operation. This is normally a comma-separated list of flag names and *option=value*\ pairs.
Logging:	Indicates whether logging is defined and supported on this volume. An undefined log type is included to support old versions of the Logical Storage Manager.
Writeback:	Indicates whether the volume is to write back on read failure. If set to Yes, an attempt will be made to fix a read error from a participating plex.
Putil0:	Permanent utility field 0. This is reserved for Logical Storage Manager use, but can be changed. The maximum length of all Putil fields is 14 characters.
Putil1:	Permanent utility field 1. This field is reserved, but can be changed.

Field	Description
Putil2:	Permanent utility field 2. This field is reserved, but can be changed.
Tutil0:	Temporary utility field 0. This is reserved for LSM use, but can be changed. The maximum length of all Tutil fields is 14 characters.
Tutil1:	Temporary utility field 1. This field is reserved, but can be changed.
Tutil2:	Temporary utility field 2. This field is reserved, but can be changed.
Kernel State:	The kernel state of this volume. These states are distinguished as follows: – Enabled — The volume device can be used. This is the default state. – Detached — The volume device cannot be used, but ioctls will still be accepted. – Disabled — The volume cannot be used for any operations.
Number of IO Failures:	The number of failed I/O operations on this volume since the last boot. This field cannot be changed.

E.2 Plex Management

The following sections provide information on menus and forms relating to plex management.

E.2.1 Plex Menus

Advanced-Ops → Plex

The Advanced-Ops menu provides access to the following plex-related menus:

- Create
- Remove Plexes
- Associate Plexes
- Dissociate Plexes
- Attach Plexes
- Detach Plexes
- Help

The Help selection accesses a Help window that displays information relevant to the plex operations.

The plex Advanced-Ops menus are described in the following list:

- Create

  Advanced-Ops → Plex → Create

  This operation creates a plex. You can select one or more subdisks to be associated with the new plex after creation.

  Form: Plex Create Form

- Remove Plexes

  Advanced-Ops → Plex → Remove plexes

  This operation removes the selected plexes. This is a permanent operation and cannot be undone. Any subdisks associated with the plex will be dissociated and left behind.

  Requirements:

  – At least one plex icon must be selected.

  – If the selected plex is associated with a volume, it must be dissociated before it can be removed.

- Associate Plexes

  Advanced-Ops → Plex → Associate Plexes

  This operation associates one or more selected plexes with the selected volume. If the volume is started, LSM begins to bring the plex up to date by copying all necessary data to the plex. This might take time.

  Requirements:

  – A volume icon and at least one plex icon must be selected.

  – Only unassociated plexes can be associated.

- Dissociate Plexes

  Advanced-Ops → Plex → Dissociate Plexes

  This operation dissociates one or more selected plexes from their parent volumes. This operation will fail if the plex cannot be dissociated. For example, the last plex in a started volume cannot be dissociated.

  Requirements:

  – At least one plex icon must be selected.

  – Only associated plexes can be dissociated.

  – Before the last plex in a volume can be dissociated, that volume must be stopped.

- Attach Plexes

  Advanced-Ops → Plex → Attach Plexes

This operation attaches one or more selected plexes to their parent volumes. A plex must be detached but still associated with an enabled volume in order to be attached; the plex is actually being reattached with its parent volume.

Requirements:

- At least one plex icon must be selected.
- A plex must be detached before it can be attached.
- Only a plex associated with an enabled volume can be attached.

• Detach Plexes

Advanced-Ops → Plex → Detach Plexes

This operation detaches one or more selected plexes from their parent volumes. A detached plex is inaccessible for reads and writes, but is still associated with the volume.

Requirements:

- At least one plex icon must be selected.
- Only associated plexes can be detached.
- This operation is not permitted when the specified plex is the last valid plex on the volume.

E.2.2 Plex Forms

Some plex operations result in the appearance of forms, which must be completed in order for that operation to proceed. Most forms provide a Help button, which contains information relevant to the fields and other aspects of that particular form.

The following forms are accessed via plex-related selections from the Advanced-Ops menu:

• Plex Create Form

Advanced-Ops → Plex → Create

The following table describes the fields in this form. Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Plex name	The name of the plex to be created. The name must be unique within this disk group. The maximum length of this field is 31 characters.
Plex state	The plex utility state. This is reserved for use by usage types. This field is optional.

Field	Description
Volume	The name of the volume that this plex should be associated with. The name must be a valid volume name in this disk group. The maximum length of this field is 31 characters. This field is optional.
Layout	The desired layout for the plex. A concatenated plex is a plex with associated subdisks that are both sequentially and contiguously arranged. A striped plex is a plex that distributes data evenly across each of its associated subdisks. The default is Concatenated.
Stripe width	The width of the stripes on the plex. The stripe width must be a number greater than 0. If no units are specified, sectors are assumed. The maximum length of this field is 14 characters. If a striped plex layout is specified, this field is required. This field must be blank if a concatenated plex layout is specified.
Subdisks	The number of subdisks associated with the plex. This field is read-only and cannot be changed.
Comment	An appropriate comment for the plex. The maximum length of the comment is 40 characters. This field is optional.
Errors	Indicates whether the plex should participate in LSM error policies. The default is Participate.
Putil0	Permanent utility field 0. This is reserved for LSM use, but may be changed. The maximum length of all Putil fields is 14 characters. This field is optional.
Putil1	Permanent utility field 1. This field is reserved, but can be changed. This field is optional.
Putil2	Permanent utility field 2. This field is reserved, but can be changed. This field is optional.

E.2.3 Plex Properties Forms

The following list describes the properties form that reveals the properties of a particular plex:

- Plex Properties Form

  To access the plex properties form, click the MB3 mouse button on the desired plex icon.

  This form provides detailed information on the attributes of a particular plex. The following table describes the fields in this form.

  The fields in this form are read/write fields, unless listed as read-only. Properties of the plex can be changed via this form by altering the current values in the appropriate read/write fields and then clicking on Apply.

Field	Description
Plex name	The name of the plex. The name must be unique within this disk group. The maximum length of this field is 31 characters. The plex name can be changed by entering another name in this field.
Plex state	The plex utility state. This is reserved for use by usage types. This field is read-only and cannot be changed.
Volume	The name of the volume that this plex should be associated with. This field is read-only and cannot be changed.
Layout	The layout of the plex: concatenated or striped. A concatenated plex is a plex with associated subdisks that are both sequentially and contiguously arranged. A striped plex is a plex that distributes data evenly across each of its associated subdisks. This field is read-only and cannot be changed.
Stripe width	The width of the stripes on the plex. If Striped plex layout has been specified, this field indicates the stripe width. This field should be blank if Concatenated plex layout has been specified. This field is read-only and cannot be changed.
Subdisks	The number of subdisks associated with the plex. This field is read-only and cannot be changed.
Log Subdisk	This field shows the name of the subdisk that is being used for logging on this plex. If there is no associated Dirty Region Logging subdisk (no logging in effect), this field is blank. This field is read-only and cannot be changed.
Comment	An appropriate comment for the plex. The maximum length of the comment is 40 characters.
Errors	Indicates whether the plex participates in LSM error policies. This field is read-only and cannot be changed.
Putil0	Permanent utility field 0. This is reserved for use, but can be changed. The maximum length of all Putil fields is 14 characters.
Putil1	Permanent utility field 1. This field is reserved, but can be changed.
Putil2	Permanent utility field 2. This field is reserved, but can be changed.
Tutil0	Temporary utility field 0. This is reserved for LSM use, but can be changed. The maximum length of all Tutil fields is 14 characters.
Tutil1	Temporary utility field 1. This field is reserved, but can be changed.
Tutil2	Temporary utility field 2. This field is reserved, but can be changed.

Field	Description
Kernel State	The accessibility of the plex. This field is read-only and cannot be changed.
Length	The length of the plex. This field is read-only and cannot be changed.
Number of I/O failures	The number of failed I/O operations on this plex since the last boot. This field is read-only and cannot be changed.

E.3 Subdisk Management

The following sections provide information on menus and forms relating to subdisk management.

E.3.1 Subdisk Menus

You access the subdisk Advanced-Ops menu as shown here:

Advanced-Ops → Subdisk

This menu provides access to the following subdisk operations:

- Create
- Remove Subdisks
- Associate Subdisks
- Associate as Log Sd
- Dissociate Subdisks
- Join Subdisks
- Split the Subdisk
- Help

The Help selection accesses a Help window that displays information relevant to the subdisk operations.

The following list describes how to access the subdisk menus:

- Create

  Advanced-Ops → Subdisk → Create

  This operation creates a subdisk on the selected LSM disk. An LSM disk must be selected.

  Form: Subdisk Create Form (described in Section E.3.2).

- Remove Subdisks

  Advanced-Ops → Subdisk → Remove Subdisks

This operation removes the selected subdisks. This is a permanent operation and cannot be undone.

Requirements:

- At least one subdisk icon must be selected.
- If the selected subdisk is associated with a plex, it must be dissociated before it can be removed. Only free subdisks can be removed.

- Associate Subdisks

  Advanced-Ops → Subdisk → Associate Subdisks

  This operation associates one or more subdisks with the selected plex.

  Requirements:

  - A plex icon and at least one subdisk icon must be selected.
  - Only unassociated (free) subdisks can be associated.

- Associate as Log Subdisk

  Advanced-Ops → Subdisk → Associate as Log Sd

  This operation associates the selected subdisk as a log subdisk with the selected plex. The resulting log subdisk icon has double borders to distinguish it from normal subdisks.

  Requirements:

  - A plex icon and a subdisk icon must be selected.
  - Only unassociated (free) subdisks can be associated.
  - The selected plex cannot already have a log subdisk.
  - Subdisks must be 2 or more sectors to enable logging in noncluster environments and 65 or more sectors for TruCluster environments.

- Dissociate Subdisks

  Advanced-Ops → Subdisk → Dissociate Subdisks

  This operation dissociates one or more selected subdisks from their parent plexes. Both log subdisks and normal subdisks can be dissociated.

  Requirements:

  - At least one subdisk icon must be selected.
  - Only associated subdisks can be dissociated.
  - The last subdisk associated with a plex that is currently associated with a volume cannot be dissociated. The plex must be dissociated from its volume first.

- Join Subdisks

  Advanced-Ops → Subdisk → Join Subdisks

This operation joins the selected subdisks to create a single subdisk. The resulting subdisk has the offset and name of the first subdisk (as arranged on the disk) and its length is the sum of the subdisk lengths.

Requirements:

– At least two subdisk icons must be selected.

– The subdisks must be contiguous on the disk.

– If the subdisks are associated, they must all be associated with the same plex and be contiguous on that plex.

– Logging subdisks and subdisks associated with striped plexes cannot be joined.

- Split a Subdisk

  Advanced-Ops → Subdisk → Split the Subdisk

  This operation splits the selected subdisk into either two or many parts. The resulting subdisks will occupy the same region on the disk that the previous subdisk occupied. If the subdisk is associated with a plex, the resulting subdisks will also be associated with that plex.

  From the Split the subdisk menu, a submenu allows you to indicate whether the subdisk is to be split into two or several parts:

  – Into 2 Subdisks — Split the selected subdisk into 2 subdisks.

  – Into More Than 2 Subdisks — Split the selected subdisk into several subdisks.

  Requirements:

  – Only one subdisk icon can be selected.

  – Logging subdisks cannot be split.

  Forms:

  – Subdisk Split Into Two

  – Subdisk Split Into Many (Section E.3.2)

E.3.2 Subdisk Forms

Some subdisk operations result in the appearance of forms, which must be completed in order for that operation to proceed. Most forms provide a Help button, which contains information relevant to the fields and other aspects of that particular form.

The following forms are accessed via subdisk-related selections from the Advanced-Ops menu:

- Subdisk Create Form

Advanced-Ops → Subdisk → Create

This form creates a subdisk according to your specifications. The following table describes the fields in this form.

Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Disk name	The name of the LSM disk on which the subdisk is to be created. This field is read-only and cannot be changed.
Subdisk name	The name of the subdisk to be created. The name must be unique within this disk group. The maximum length of this field is 31 characters.
Disk offset	The length into the disk where this subdisk should be located. If no units are specified, sectors are assumed. This offset should not place this subdisk within the bounds of another subdisk on the disk or past the end of the disk. Only valid positive numbers are allowed in this field.
Subdisk length	The length of the subdisk to be created. If no units are specified, sectors are assumed. The length should not place this subdisk within the bounds of another subdisk on the disk or past the end of the disk. Only valid positive numbers are allowed in this field.
Plex name	The name of the plex with which the subdisk is to be associated. This must be a valid plex that already exists in this disk group. The maximum length of this field is 31 characters. This field is optional.
Plex offset	The offset of this subdisk into its associated plex. Only valid positive numbers are allowed in this field. This field is required only if a plex has been specified for association. If the subdisk is not to be associated with a plex, this field must be left blank.
Comment	An appropriate comment for the subdisk. The maximum length of the comment is 40 characters. This field is optional.
Putil0	Permanent utility field 0. This is reserved for Logical Storage Manager use, but may be changed. The maximum length of all Putil fields is 14 characters. This field is optional.
Putil1	Permanent utility field 1. This field is reserved, but may be changed. The maximum length of this field is 14 characters. This field is optional.
Putil2	Permanent utility field 2. This field is reserved, but may be changed. The maximum length of this field is 14 characters. This field is optional.

- Subdisk Split Into Two

Advanced-Ops → Subdisk → Split the Subdisk –> Into 2 Subdisks

This form splits the selected subdisk into exactly two subdisks. The first subdisk retains the name and size of the original one; the second subdisk adopts the name and size specified in this form. The following table describes the fields in this form.

Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Present size	The size of the subdisk to be split. This field is read-only and cannot be changed.
Name of new subdisk	The name of the subdisk to be created from the original one. This must be a valid name and must be unique in this disk group.
Size of new subdisk	The size of the subdisk to be created from the original one. This must be a valid number, greater than zero. The new subdisk size must be at least one sector less than the present subdisk size.

- Subdisk Split Into Many

  Advanced-Ops → Subdisk → Split the Subdisk → Into More Than 2 Subdisks

  This form splits the selected subdisk into several subdisks of equal sizes. The first subdisk retains the name and size of the original one; the additional subdisks are automatically named by LSM. The following table describes the fields in this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Original subdisk	The name of the selected subdisk. This field is read-only and cannot be changed.
Present size	The size of the subdisk to be split. The original subdisk must contain enough sectors to accommodate the desired total number of subdisks for the split. This field is read-only and cannot be changed.
Number of new subdisks	The total number of subdisks to be created by the split. There must be a sufficient number of sectors in the original subdisk to accommodate this number. This number should be at least 2.

Requirements: The number of subdisks is limited by the amount of space left in the configuration database.

E.3.3 Subdisk Properties Forms

The following is the properties form that reveals the properties of a particular subdisk:

- Subdisk Properties Form

  To access the Subdisk Properties form, click MB3 on the desired subdisk icon. If the subdisk is undergoing analysis, press Shift–MB3 instead. This form provides detailed information on the attributes of a particular subdisk. The following table describes the fields in this form.

  The fields in this form are read/write fields, unless listed as read-only. Properties of the subdisk can be changed via this form by altering the current values in the appropriate read/write fields and then clicking on Apply.

Field	Description
Disk name	The name of the disk where the subdisk resides. This field is read-only and cannot be changed.
Subdisk name	The name of the subdisk. The name must be unique within this disk group. The maximum length of this field is 31 characters. The subdisk name can be changed by entering another name in this field.
Disk offset	The length into the disk where this subdisk is located, in sectors. This field is read-only and cannot be changed.
Subdisk length	The length of the subdisk. If no units are specified the number is assumed to be in sectors. This offset should not place this subdisk within the bounds of another subdisk on the disk or past the end of the disk. Only valid positive numbers are allowed in this field.
Plex name	The name of the plex with which the subdisk is associated. This field is read-only and cannot be changed.
Plex offset	The offset of this subdisk into its associated plex. If the subdisk is not associated, this field contains a zero. This field is read-only and cannot be changed.
Comment	An appropriate comment for the subdisk. The maximum length of the comment is 40 characters.
Log Subdisk	Indicates whether this subdisk is a Dirty Region Logging subdisk. This field is read-only and cannot be changed.
Putil0	Permanent utility field 0. This is reserved for LSM use, but may be changed. The maximum length of all Putil fields is 14 characters.
Putil1	Permanent utility field 1. This field is reserved, but may be changed. The maximum length of this field is 14 characters.

Field	Description
Putil2	Permanent utility field 2. This field is reserved, but may be changed. The maximum length of this field is 14 characters.
Tutil0	Temporary utility field 0. This is reserved for LSM use, but may be changed. The maximum length of all Tutil fields is 14 characters.
Tutil1	Temporary utility field 1. This field is reserved, but may be changed. The maximum length of this field is 14 characters.
Tutil2	Temporary utility field 2. This field is reserved, but may be changed. The maximum length of this field is 14 characters.
Number of IO failures	The number of failed I/O operations on this subdisk since the last boot. This field is read-only and cannot be changed.

E.4 Disk Management

The following sections provide information on menus and forms relating to disk management.

E.4.1 Disk Menus

Both the Basic-Ops and Advanced-Ops menus provide access to disk-related operations. Most menus provide a Help selection, which contains information relevant to the items and operations listed in that particular menu.

E.4.1.1 Basic-Ops Menu

You access the Basic-Ops menu by selecting:

Basic-Ops → Disk Operations

This menu provides access to disk operations involving general disk maintenance. These operations use the automated approach to disk management.

The Disk Operations menu provides the following selections:

- Add Disks
- Evacuate Subdisks
- Replace Disks
- Remove Disks
- Help

The Help selection accesses a Help window, which displays information relevant to the basic disk operations.

The following list describes the menu selections you can access via the Basic-Ops menu:

- Add Disks

  Basic-Ops → Disk Operations → Add Disks

  This operation adds a disk to the Logical Storage Manager, placing it under LSM control. This involves initializing, analyzing, and partitioning the raw disk; initializing the disk for LSM use; and adding the disk to a disk group (if requested).

  Form: Add Disks Form (described in the Disk Forms section).

- Evacuate Disks

  Basic-Ops → Disk Operations → Evacuate Disks

  This operation moves all subdisks from the selected disk to another disk in the same disk group.

  Requirements: The disk from which subdisks are to be evacuated must be selected. Both disks must belong to the same disk group.

  Forms: Evacuate Subdisks Form (described in the Disk Forms section).

- Replace Disks

  Basic-Ops → Disk Operations → Replace Disks

  This operation replaces a disk. This is normally done when a failed disk needs to be replaced with a new one. This involves initializing and partitioning the raw disk; initializing the disk for LSM use; and replacing the old disk and associated disk media records with the new disk and its information.

  Requirements: A disk icon representing a failed disk must be selected.

  Forms: Replace Disks Form (described in the Disk Forms section).

- Remove Disks

  Basic-Ops → Disk Operations → Remove Disks

  This operation removes a disk from a disk group and then removes the disk from LSM control.

  Requirements: A disk icon must be selected.

E.4.1.2 Advanced-Ops Menu

You access the Advanced-Ops menu selections by selecting:

Advanced-Ops → Disk

This menu provides access to assorted disk operations using the manual approach to disk management.

The Disk menu provides the following selections:

- Initialize
- Define
- Remove
- Online
- Offline
- Help

The Help selection accesses a Help window, which displays information relevant to the advanced disk operations.

The following list describes the menu selections you can access via the Advanced-Ops menu:

- Initialize

  Advanced-Ops → Disk → Initialize

  This operation identifies a disk to LSM and initializes the disk for LSM use. This involves installing a disk header and writing an empty configuration on the disk. A disk access record is created for the disk, unless such a record already exists.

  Requirement: The disk should not already be initialized.

  Form: Disk Init Form (described in the Disk Forms section).

- Define

  Advanced-Ops → Disk → Define

  This operation defines a disk access record that enables LSM to scan the disk. This makes the disk accessible, but does not initialize the disk.

  Form: Define Disk Form (described in the Disk Forms section).

- Remove

  Advanced-Ops → Disk → Remove

  This operation removes the LSM disk associated with the selected partitions from LSM control by removing the associated disk access records. If all partitions on a given disk are selected for removal at once, the disk is effectively removed from LSM control.

  Requirements:

  – At least one partition icon corresponding to an LSM disk must be selected.

  – The LSM disks corresponding to the selected partitions cannot belong to a disk group at the time of removal.

- Online

  Advanced-Ops → Disk → Online

  This operation places the disk access record on a specified partition in an online state. During searches for disk IDs or members of a disk group, online disks are checked.

  Form: Disk Online Form (described in the Disk Forms section).

- Offline

  Advanced-Ops → Disk → Offline

  This operation places the disk access record on the selected partitions in an offline state. During searches for disk IDs or members of a disk group, offline disks are ignored.

  Requirements:

  – At least one partition icon must be selected.

  – The disks corresponding to the selected partitions must be initialized.

  – The selected partition icon cannot be in use (shaded and associated with an LSM disk).

E.4.2 Disk Forms

Some disk operations result in the appearance of forms. You must complete these forms in order for that operation to proceed. Most forms provide a Help button that provides access to information relevant to the fields and other aspects of that form.

E.4.2.1 Basic-Ops Forms

The following forms are accessed via disk-related selections from the Basic-Ops menu.

- Add Disks Form

  Basic-Ops → Disk Operations → Add Disks

  This form places a disk under Logical Storage Manager control. The following table describes the fields in this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
New disk name	The name of the new physical disk in the form `dsk`*nn*, for example, `dsk10`. The name must be unique within this disk group. You can also place specific partitions on a disk under LSM control. For example, `dsk3g` would put the `g` partition on `dsk3` under LSM control.
Disk group	The name of the disk group to which this disk is to be added. The named disk group must exist. If no name is provided, it will not be added to a disk group. This field is optional.

- Replace Disks Form

  Basic-Ops → Disk Operations → Replace Disks

  This form replaces an existing LSM disk that has failed with another one. The following table describes the fields in this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Old LSM disk name	The name of the failed (collapsed or disconnected) LSM disk in this disk group. This field is read-only and cannot be changed.
New physical disk name	The name of the new physical disk that is to replace the existing one. The name should be in the form `dsk`*nn*, for example, `dsk10`. The new name must be unique in this disk group.

- Evacuate Subdisks Form

  Basic-Ops → Disk Operations → Evacuate Subdisks

  This form transfers subdisks from one LSM disk to another. The following table describes the fields in this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Disk group name	The name of the disk group to which both disks belong. Both disks must share the same disk group.
Evacuate From	The name of the LSM disk from which the subdisks are to be evacuated.
To	The name of the LSM disk to which the subdisks are to be moved. This field is optional. However, if no target disk is specified, the subdisks are evacuated to one or more random disks (depending on disk space availability).

E.4.2.2 Advanced-Ops Forms

The following forms are accessed via disk-related selections from the Advanced-Ops menu:

- Disk Init Form

  Advanced-Ops → Disk → Initialize

  This form initializes a disk for LSM use. The following table describes the fields in this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Public Device	The pathname of the device node that represents a partition available for use. This name must be a valid entry in `/dev/disk`. A name in the form `dsk`*nn* is used to assign the full disk under LSM control. The disk `dsk`*nn* would be added as a sliced LSM disk. Before a sliced disk can be defined, change the disk label to have LSM disk label tags. A name in the form `dsk`*nnp* is used to assign partition *p* on disk `dsk`*nn* under LSM control. The disk partition `dsk`*nnp* is added as a simple LSM disk.
Device Type	The desired disk type. The simple type (default) assumes that the public and private regions are stored on the same disk partition, with the public region following the private region. The sliced type assumes that the public and private regions are stored on different disk partitions. Before initializing the disk, change the disk label to have LSM disk label tags. The `nopriv` type has no private region and log and configuration copies cannot be written to the disk.
Public length (0 for whole device)	The length of the public region of the disk. If zero is provided as the length, the Logical Storage Manager computes a default value from available partition table information. This length must be valid and cannot exceed the length of the disk.
Private Length	The length of the private region of the disk. When one is not specified, LSM chooses a default value. This length must be valid and cannot exceed the length of the disk. For a sliced disk, the length cannot exceed the size of the partition chosen for the private region. This field is optional.
Number of config copies	The number of configuration copies to be stored in the private region of this disk. The default value is one copy.
Comment:	A comment appropriate for the LSM disk. The maximum length of the comment is 40 characters. This field is optional.

- Define Disk Form

Advanced-Ops → Disk → Define

This form defines a disk. The following table describes the fields in this form.

Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Public Device	The pathname of the device node that represents a partition available for use. This name must be a valid entry in `/dev/disk`. A name in the form `dsk`*nn* is used to assign the full disk under LSM control. The disk `dsk`*nn* would be added as a sliced LSM disk. A name in the form `dsk`*nnp* is used to assign partition *p* on disk `dsk`*nn* under LSM control. The disk partition `dsk`*nnp* would be added as a simple LSM disk.
Device Type	The desired disk type. The simple type (default) assumes that the public and private regions are stored on the same disk partition, with the public region following the private region. The sliced type assumes that the public and private regions are stored on different disk partitions. The nopriv type has no private region and log and configuration copies cannot be written to the disk.
Public Length (0 for whole disk)	The length of the public region of the disk. If zero is provided as the length, LSM computes a default value from available partition table information. This length must be valid and cannot exceed the length of the disk.
Offline	Indicates whether to initially place the disk in the offline state. The default is No.
Comment	A comment appropriate for this Logical Storage Manager disk. The maximum length of the comment is 40 characters. This field is optional.

- Disk Online Form

  Advanced-Ops → Disk → Online

  This form places a disk on line. The following table describes the fields in this form.

Field	Description
Device name	The disk access name of the disk to be placed online. This must be a valid disk access name. This field is required.

- Free Space Form

  To access the free space form, click MB3 on a gap between subdisk icons in an LSM disk icon.

This form provides information about a specific region of an LSM disk that contains free space.

Free space results when subdisks are removed for some reason, making the space that they occupied available for use. Free space is visually represented as a gap or hole between subdisks that reside on an LSM disk icon. The following table describes the fields in the form. All fields in this form are read-only and cannot be changed.

Field	Description
Device	The name of the LSM disk where this free space resides.
Hole offset	The offset into the LSM disk where this free space extent begins.
Hole size	The size of this free space extent. Specify the units used from the Options pull down menu.

E.4.2.3 Disk Properties Forms

Properties forms exist for LSM disks, physical disks, and partitions. The following list describes these forms:

- LSM Disk Properties Form

  To access the LSM disk properties form, click MB3 on desired LSM disk icon. (If the LSM disk icon is undergoing analysis, press Shift–MB3 instead.)

  This form provides detailed information on the attributes of a particular LSM disk that is under LSM control. The information displayed in this form actually corresponds to the disk media record associated with a disk. The following table describes the fields in this form.

  The fields in this form are read/write fields, unless listed as read-only. Properties of the disk can be changed via this form by altering the current values in the appropriate read/write fields and then clicking on Apply.

Field	Description
LSM disk name	The name of the LSM disk.
Disk Access	The name of the disk access record that corresponds to this disk media record. This field is read-only and cannot be changed.
Disk Type	The type with which this disk media record was created. This field is read-only and cannot be changed.
Public Region	The name of the public region of this disk. This field is read-only and cannot be changed.

Field	Description
Private Region	The name of the private region of this disk. If there is no private region then this field will be blank. This field is read-only and cannot be changed.
Public Region Offset	The offset, in sectors, of the public region on the disk. This field is read-only and cannot be changed.
Private Region Offset	The offset, in sectors, of the private region on the disk. If there is no public region, then this field will display zero. This field is read-only and cannot be changed.
Public Region Length	The length, in sectors, of the public region on the disk. This field is read-only and cannot be changed.
Private Region Length	The length, in sectors, of the private region on the disk. If there is no private region, this field will display zero. This field is read-only and cannot be changed.
Disk Attributes	The attributes of this LSM disk. This field is read-only and cannot be changed.
Comment	The user-specified comment for this LSM disk. The maximum length of the comment is 40 characters.
Putil0	Permanent utility field 0. This is reserved for LSM use, but may be changed. The maximum length of all Putil fields is 14 characters.
Putil1	Permanent utility field 1. This field is reserved, but may be changed.
Putil2	Permanent utility field 2. This field is reserved, but may be changed.
Tutil0	Temporary utility field 0. This field is reserved, but may be changed. The maximum length of all Tutil fields is 14 characters.
Tutil1	Temporary utility field 1. This field is reserved, but may be changed. The maximum length of this field is 14 characters.
Tutil2	Temporary utility field 2. This field is reserved, but may be changed. The maximum length of this field is 14 characters.
Maximum Free Space	The maximum amount of free space available on this LSM disk. This does not take disk extents into account. This number assumes every free sector on the LSM disk is usable. This field is read-only and cannot be changed.

- Physical Disk Properties Form

  To access the physical disk properties form, click MB3 on the desired physical disk icon.

  This form provides detailed information on the attributes of a particular physical disk. The following table describes the fields in this form.

All fields in this form are read-only and cannot be changed.

Field	Description
Device	The raw device node for this physical disk.
Device Type	A brief description the device type. Possible device types include SCSI hard drive and Floppy.
Cylinders	The number of cylinders on this disk.
Tracks	The number of tracks per cylinder.
Sectors	The number of sectors per track.
Sector Size	The size, in bytes, of each sector on this disk.
Total Size	The total size of the disk, in sectors.

- Partition Properties Form

  To access the partition properties form, click MB3 on the desired partition icon.

  This form provides detailed information on the attributes of a particular partition. The following table describes the fields in this form.

  All fields in this form are read-only and cannot be changed.

Field	Description
Device	The device node that the LSM Visual Administrator uses to communicate with this disk.
Start Sector	The sector on the physical disk where this partition begins.
Size	The length of this partition.
Type	The identification tag associated with this partition.
Disk Media	The disk media record that corresponds to this partition. If this field is empty, the partition has not been initialized with a disk media record.

E.5 Disk Group Management

The following sections provide information on menus and forms relating to disk group management.

Note

With the Visual Administrator, partition icons represent partitions described by disk access records.

E.5.1 Disk Group Menus

You access disk group operations via the Advanced-Ops menu, as shown here:

Advanced-Ops → Disk Group

The Advanced-Ops menu provides access to the following disk-related menus.

- Initialize
- Import Disk Groups
- Deport Disk Groups
- Add Disk
- Remove Disks
- Disconnect Disks
- Reconnect Disks
- Help

The Help selection accesses a Help window that displays information relevant to the disk group operations.

The following list describes the disk group menu options:

- Initialize

  Advanced-Ops → Disk Group → Initialize

  This operation defines a new disk group with a name you specify. The new disk group contains one or more LSM disks corresponding to the partitions you select.

  Requirements: At least one partition icon must be selected.

  Form: Initialize Disk Group Form (described in the Disk Group Forms section).

- Import Disk Group

  Advanced-Ops → Disk Group → Import Disk Group

  This operation imports a disk group to make that disk group available on the local machine. If the name of a deported disk group is known, this operation can be used to make that disk group accessible again.

  Form: Import Disk Group Form (described in the Disk Group Forms section).

- Deport Disk Group

  Advanced-Ops → Disk Group → Deport Disk Group

This operation disables access to a disk group. A deported disk group is no longer accessible and its view window disappears. Once deported, a disk group can be reimported.

Requirements: A disk group cannot be deported if any volumes in that disk group are currently open.

Form: Deport Disk Group Form (described in the Disk Group Forms section).

- Add Disk

  Advanced-Ops → Disk Group → Add Disk

  This operation adds an LSM disk corresponding to the selected partition icon to a disk group. This involves creating a disk media record for the disk to be added. Partitions representing disks that already belong to disk groups cannot be added to disk groups.

  Requirements:

  – One partition icon must be selected.

  – The selected partition cannot already belong to a disk group.

  – Only one disk can be added to a disk group at a time.

  Form: Add Disk Form (described in the Disk Group Forms section).

- Remove Disks

  Advanced-Ops → Disk Group → Remove Disks

  This operation removes the selected LSM disks from a disk group. Disks are removed from the disk group in which they reside. Any subdisks that exist on the selected disks must be removed before the disk can be removed.

  Requirements:

  – At least one LSM disk icon must be selected.

  – Only disks associated with the specified disk group can be removed.

  – Disks containing any subdisks cannot be removed.

  – Only disks in the same disk group can be selected for removal in a single operation.

  – The last disk in a disk group cannot be removed. The disk group itself must be deported in order for its last disk to be removed.

- Disconnect Disks

  Advanced-Ops → Disk Group → Disconnect Disks

This operation disables the selected LSM disk, making it unavailable for use within its disk group. This involves dissociating the disk media record from its disk access record.

Requirements:

– At least one LSM disk icon must be selected.

– The LSM disk icons must contain a disk media record at the time of selection.

- Reconnect Disks

  Advanced-Ops → Disk Group → Reconnect Disks

  This operation enables an LSM disk that has previously been disconnected. This involves connecting the selected LSM disk's disk media record with the selected disk access record. Although the LSM disk must be disconnected, it does not necessarily have to be reconnected to its former partition (disk access record).

  Requirements:

  – One LSM disk icon and one partition icon must be selected.

  – Neither the LSM disk icon nor the partition icon can already be connected.

E.5.2 Disk Group Forms

Some disk group operations result in the appearance of forms, which must be completed in order for that operation to proceed. Most forms provide a Help button, which contains information relevant to the field and other aspects of that particular form.

The following forms are accessed via disk group-related selections from the Advanced-Ops menu:

- Initialize Disk Group Form

  Advanced-Ops → Disk Group → Initialize

  This form defines a new disk group consisting of selected disks.

  The following table describes the field in this form.

Field	Description
Disk group	The name of the new disk group. This must be a valid and unique name. This field is required. This is a read/write field.

- Import Disk Group Form

  Advanced-Ops → Disk Group → Import Disk Group

This form makes the specified disk group available to the system. The following table describes the field in this form.

Field	Description
Disk group	The name of the disk group to be imported and made available to the system. This must be a valid and unique disk group name. This field is required. This is a read/write field.

- Deport Disk Group Form

  Advanced-Ops → Disk Group → Deport Disk Group

  This form makes the specified disk group inaccessible to the system. The following table describes the fields in this form.

Field	Description
Disk group	The name of the disk group to be deported and made inaccessible to the system. This must be a valid disk group.

  Requirements: The root disk group (rootdg) cannot be deported.

- Add Disk Form

  Advanced-Ops → Disk Group → Add Disk

  This form adds an LSM disk to a disk group. The following table describes the fields in this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Disk group	The name of the disk group to which the LSM disk is to be added. This must be a valid disk group. This field is required.
Disk media name	The name of the LSM disk to be created. The disk media name must be unique. By default, a unique name is generated. If this field is left blank, then the disk access name is used.

E.6 Projection Analysis

The following sections provide information on menus and forms relating to projection and analysis. In addition, tables are provided to summarize various aspects of projection and analysis behavior. You can access these operations as follows:

- Projection

  Projection operations are accessed via the Projection menu. This menu is located in view windows such as View of the rootdg disk group. The

Projection menu starts or stops projection, and highlights any free subdisk icons.

Projection can also be started or stopped by pressing Shift–MB2 with the pointer positioned on the desired icon.

- Analysis

  Analysis operations are accessed via the Analyze menu. This menu is located in view windows such as View of the rootdg disk group. The Analyze menu can be used to starts or stops analysis and sets analysis-related preferences.

E.6.1 Projection

Icon projection provides you with visual information about the relationships between icons. When projection is started for an icon, all other icons (representing LSM objects) associated with that particular one are highlighted, no matter which views they occupy. Icons can be placed under projection either individually or in multiples. Projection highlighting can accumulate on a given icon when that icon is undergoing projection from more than one source.

E.6.2 Projection Menus

The following list describes the menus, submenus, and menu selections you can access via the Projection menu:

- Icon Projection

  Projection → Icon Projection

  This menu provides access to projection options used to start or stop projection for icons.

- Start

  Projection → Icon Projection → Start

  This option starts projection for the selected icons. When projection is started, all icons related to the selected icons are highlighted. Highlighting occurs for related icons in any view windows. If the selected icon has no associated objects, the Visual Administrator issues a warning to this effect.

  Requirements:

  – At least one icon must be selected.

  – Physical disk and partition icons cannot be selected for projection.

  – The selected icons must be associated with at least one other icon in order for projection to take effect.

- Stop

  Projection → Icon Projection → Stop

  This option stops projection for the selected icons. When projection is stopped, all icons related to the selected icons lose their projection highlighting.

  Requirement: At least one icon must be selected. If the selected icon is not undergoing projection, the Visual Administrator ignores the stop request.

- Stop All

  Projection → Icon Projection → Stop All

  This option stops projection for all icons that are currently undergoing selection.

- Show Free Subdisks

  Projection → Show Free Subdisks

  This menu selection determines whether free subdisks should be highlighted or not. When Show Free Subdisks is turned on, the Visual Administrator highlights all unassociated subdisks (representing unallocated disk space). Once turned on, any future free subdisks are automatically highlighted. Free subdisk icons can be used by designating them to objects, but the LSM Visual Administrator interface cannot automatically use free subdisks as free space. Free subdisk projection is either started or stopped across all Visual Administrator views. The start or stop preference is also retained for a particular user in future sessions.

  From the Show Free Subdisks menu, a submenu allows you to indicate whether or not to highlight free subdisks:

Option	Description
Start	Start highlighting free subdisks immediately and continue to do so until instructed to stop.
Stop	Stop highlighting free subdisks.

E.6.3 Projection Relationships

Table E–3 summarizes the projection relationships that are highlighted for particular icon types. If no icons of the correct type are associated with the selected icon, then nothing is highlighted.

Table E–3: Projection Table

Icon Selected	Icons Highlighted
Volume	All subdisks associated with any plex associated with the volume
Plex	All subdisks associated with the plex
Subdisk	Associated plex and volume, and all other subdisks associated with the plex
LSM Disk	All plexes associated with the subdisks that reside on the disk

E.7 Analysis

Analysis is the LSM Visual Administrator's way of displaying statistics on the performance of various LSM objects.

Statistics are displayed both visually (via color or pattern) and numerically (via pop-up statistics forms).

E.7.1 Analysis Menus

The following menu selections are accessed via the Analyze menu:

- Start

  Analyze → Start

  This menu selection begins analysis of the selected icons. These icons are added to the list of objects being analyzed. Only volume and LSM disk icons can be analyzed. Once analysis is activated, the selected icons begin to display information about their performance characteristics.

  Requirement: At least one volume or LSM disk icon must be selected.

- Stop

  Analyze → Stop

  This menu selection terminates analysis of the selected icons. These icons are removed from the list of objects being analyzed. When analysis stops, the selected icons return to their preanalysis states. When analysis is stopped for one icon, other icons undergoing analysis are not affected.

  Requirements:

  – At least one volume or LSM disk icon must be selected.

  – The selected icons must be undergoing analysis.

- Stop All

  Analyze → Stop All

This menu selection automatically terminates analysis of all icons in all views. All icons return to their preanalysis states.

Requirements: Analysis must be in effect.

- Parameters

Analyze → Parameters

This menu selection accesses the Analysis Parameters form, which sets user preferences for how analysis is to be conducted.

Form: Analysis Parameters Form (described in the Analysis Forms section).

E.7.2 Analysis Forms

The following forms are accessed via the Analyze menu:

- Analysis Parameters Form

Analyze → Parameters

This form sets user preferences for conducting analysis. The following table describes the fields in this form.

Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Sample Rate	Determines the time interval between data samples. This field is divided into two sections: the slider bar selects the interval (1–60) and the menu to the right selects units of time (seconds or minutes). The default is 5 seconds. A shorter interval means the data will be updated more often, but is also a higher load on the system.
Volume Parameters	Specifies the high and low values that decide the coloring (or pattern) of the volume icons.
Disk Parameters	Specifies the high and low values that decide the coloring (or pattern) of the LSM disk icons.
Subdisk Parameters	Specifies the high and low values that decide the coloring (or pattern) of the subdisk icons.
Log File	The name of the file for the statistics log. If the file does not exist, it is binary file created. The file name is taken to be relative unless a path name is given. To stop logging to the file, delete the file name text in this field. This field is optional. To view the log file, you must run `/usr/bin/lsmlog2text` *filename* on the file to process it for viewing.

Requirements:

- For each set of high/low parameters, the high parameter must be greater than the low parameter.
- You must have access to the specified log file.

- Analysis Statistics Form

  To access the analysis statistics form, click MB3 on desired the icon that is being analyzed.

  This form displays analysis statistics relevant to the selected volume or LSM disk icon. This form applies only to volume or disk icons that are undergoing analysis. The following table describes the fields in this form. All fields in this form are read-only and cannot be changed.

Field	Description
Reads	The number of times the object was read from during the last interval.
Writes	The number of times the object was written to during the last interval.
Total R/W	The total number of reads and writes during the last interval.
Blocks Read	The number of disk blocks read from the object during the last interval.
Blocks Written	The number of disk blocks written to the object during the last interval.
Total Blocks	The total number of blocks read from or written to the object during the last interval.
Avg Read Time	The average time, in milliseconds, that it took for a read operation to complete. This is equal to the number of number of reads during the last interval divided by the total time spent on reads.
Avg Write Time	The average time, in milliseconds, that it took for a write operation to complete. This is equal to the number of writes during the last interval divided by the total time spent on writes.
Interval	The actual time, in seconds, since the last data was sampled. This might vary slightly from the specified interval time due to uncontrollable variances from system to system.

  Requirement: The icon selected by clicking MB3 must be undergoing analysis.

E.7.3 Analysis Table

Table E–4 summarizes the default colors and patterns associated with the various levels of analysis. These defaults can be changed using the X resources for `dxlsm`. See the `dxlsm`(8X) reference page for more information.

Table E–4: Analysis Table

Analysis Level	Color	Bitmap Pattern
low	green	cross_weave
medium	yellow	root_weave
high	red	wide_weave

E.8 UFS Management

This section provides information on Visual Administrator UFS file system operations. You access UFS operations via the Basic-Ops menu. This menu is located in view windows, such as View of the rootdg disk group. This menu provides access to UFS operations involving general file system maintenance, and is accessed by selecting:

Basic-Ops → UFS Operations

You can access the following menu selections via the Basic-Ops menu.

- Create
- Make File System
- Mount
- Unmount
- Check File System (fsck)
- Display Properties
- Help

The Help selection accesses a Help window which displays information relevant to the available file system operations.

E.8.1 File System Menus

The following list describes the file system operations menu items:

- Create

  Basic-Ops → UFS Operations → Create

  This operation creates a file system on an underlying volume. This is done by creating a volume on one or more disks and then creating the file system on that volume.

  You can select one or more disks on which to create the volume (providing that there is sufficient space on the disks). If you do not specify any disks, LSM automatically determines which disks to use based on available free space.

From the Create menu, select the type of volume to be created from a submenu listing two of the basic types of volumes:

Type	Description
Simple	Creates a simple, concatenated volume whose subdisks are arranged both sequentially and contiguously within a plex.
Striped	Creates a volume with data spread fairly evenly across multiple disks by way of striping. Stripes are relatively small, equally sized fragments that are allocated alternately to the subdisks of each plex.

If a mirrored volume is desired, a simple or striped volume must be created and then mirrored using the Add Mirror option from the Volume Operations menu.

Requirements:

- Only disks in the same disk group can be selected.
- Only LSM disks (disks under LSM control) can be selected.
- If striping is to be in effect, at least two disks are required in order for the operation to succeed.

Forms:

- Simple Volume/FS Create
- Striped Volume/FS Create

• Make File System

Basic-Ops → UFS Operations → Make File System

This operation makes a file system on an existing volume. Select the volume on which to place the new file system, and specify the mount point if the file system is to be mounted immediately.

Requirements:

- A volume icon must be selected.
- The selected volume must be enabled.
- Only one mounted file system can exist on each volume.

Form: Make File System Form (described in Section E.8.2).

• Mount

Basic-Ops → UFS Operations → Mount

This operation mounts the file system that resides on the selected volume. This operation assumes that the selected volume already contains a valid file system. The Visual Administrator has no way of knowing whether a valid, unmounted file system already exists on a given volume. You must make sure of the existence of an unmounted file system on a volume, as well as that file system's type.

Requirements:

– A volume icon must be selected.

– A valid, unmounted file system must already exist on the selected volume.

Form: Mount File System Form (described in Section E.8.2).

- Unmount

  Basic-Ops → UFS Operations → Unmount

  This operation unmounts the file systems that resides on the selected volumes. The file system can be unmounted only if the mount point is not busy.

  Requirements:

  – At least one volume icon must be selected.

  – The selected volume must contain a mounted file system.

- Check File System

  Basic-Ops → UFS Operations → Check File System (fsck)

  This operation checks the file systems on the selected volumes for consistency (using `fsck`). The file system to be checked must currently be unmounted.

  Requirements:

  – At least one volume icon must be selected.

  – The selected volumes must contain an unmounted file system.

  Form: File System Check Form (described in Section E.8.2).

- Display Properties

  Basic-Ops → UFS Operations → Display Properties

  Display information for file systems mounted on the system. You can select the file system for which information is to be displayed from a list of all mounted file systems. If a volume is selected, the properties for the file system that resides on that volume is displayed by default.

E.8.2 File System Forms

Some file system operations result in the appearance of forms that you must complete in order for that operation to proceed. Most forms provide a Help button that provides access to information relevant to the fields and other aspects of that particular form.

E.8.2.1 Basic-Ops Forms

The following list describes how to access forms via file system-related selections from the Basic-Ops menu:

- Simple Volume/FS Create Form

  Basic-Ops → UFS Operations → Create → Simple

  This form creates a concatenated volume and then creates a file system on the new volume. The form is divided into two sections, one for volume creation and the other for file system creation. Most of the form fields are already set to the defaults for the creation of a new volume and file system. The following table describes the fields for this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Volume name	The name of the volume to be created. The name must be unique within this disk group. The maximum length of this field is 31 characters.
Volume size	The desired volume size. The size should be entered as a number followed immediately by the letter `k`, `m`, or *s* to indicate kilobytes, megabytes, or sectors, respectively. If no unit is specified, the default is sectors. The volume size should be less than or equal to the available free space of the disks.
Usage Type	The desired usage type. The `fsgen` file system is the generic usage type, which assumes that the volume is being used by a file system. The `gen` file system is the generic usage type, which makes no assumptions regarding the data content of the volume. The default is `fsgen`.
Create file system	Indicates whether a file system is to be created. When this form is invoked from the UFS Operations menu, the default is to create a file system (Yes). All fields below this field are accessible only when Yes is specified here.
FS type	UFS is the only currently supported file system type.
Mount file system	Indicates whether the file system should be mounted after creation. If the answer is Yes (the default), a mount point must also be specified in the next field. All fields below this field are accessible only when Yes is specified here.

Field	Description
Mount point	The desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator automatically creates it. This field is required if the file system is to be mounted.
Mount automatically	Indicates whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/fstab`). The default is Yes.

- Striped Volume/FS Create Form

  Basic-Ops → UFS Operations → Create → Striped

  This form creates a striped volume and creates a file system on the new volume. The form is divided into two sections, one for volume creation and the other for file system creation. Most of the form fields are already set to the defaults for the creation of a new volume. The following table describes the fields for this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Volume name	The name of the volume to be created. The name must be unique within this disk group. The maximum length of this field is 31 characters.
Volume size	The desired volume size. The size should be entered as a number followed immediately by the letter `k`, `m`, or *s* to indicate kilobytes, megabytes, or sectors, respectively. If no unit is specified, the default is sectors. If the size is not wholly divisible by the stripe width, LSM adjusts the volume size up to the next even multiple in order to create the volume. For a striped volume, the volume size should be calculated as follows: *vol_size = stripe_width * number_of_stripes * n*, where *n* is a number greater than zero. The volume size should be less than or equal to the available free space of the disks.
Usage Type	The desired usage type. The `fsgen` type is the file system generic usage type, which assumes that the volume is being used by a file system. The `gen` type is the generic usage type, which makes no assumptions regarding the data content of the volume. The default is `fsgen`.
Number of stripes	The number of stripes that the volume's plex is to have. This is effectively the number of disks on which the volume is to be created. If some number of disks have already been selected, that number of stripes appears in this field. This number corresponds to the number of disks across which data will be striped. If no number is specified, an appropriate number (usually 2) is used.

Field	Description
Stripe width	The width of the stripes on the plex that this volume will have. The value specified might be optimized for the particular application. However, the default value of 128 sectors is as a good stripe width for most systems.
Create file system	Indicates whether a file system is to be created. When this form is invoked from the UFS Operations menu, the default is to create a file system (Yes). All fields below this field are accessible only when Yes is specified here.
FS type	UFS is the only currently supported file system type.
Mount file system	Indicates whether the file system should be mounted after creation. If the answer is Yes (the default), a mount point must also be specified in the next field. All fields below this field are accessible only when Yes is specified here.
Mount point	The desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator automatically creates it. This field is required if the file system is to be mounted.
Mount automatically	Indicates whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/fstab`). The default is Yes.

- Make File System Form

Basic-Ops → UFS Operations → Make

This form makes a file system (using `newfs`) according to your specifications. The following table describes the fields for this form.

Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless they are listed as read-only.

Field	Description
Device name	Displays the block device on which to make the file system, which corresponds to the name of the selected volume. This field is read-only and cannot be changed.
File system size	Displays the length of the file system to be made. If no units are specified, sectors are assumed. This length should typically correspond to the length of the volume on which the file system is to be made, although it can be altered for special circumstances.
FS Type	UFS is the only currently supported file system type.
Mount file system	Indicates whether the file system should be mounted after creation. If the answer is Yes (the default), a mount point must also be specified in the next field. All fields below this field are accessible only when Yes is specified here.

Field	Description
Mount point	The desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator automatically creates it. This field is required if the file system is to be mounted.
Mount automatically	Indicates whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/fstab`). Yes is the default.

- Mount File System Form

  Basic-Ops → UFS Operations → Mount

  This form mounts a file system that already exists on a selected volume. The following table describes the fields for this form.

  Most fields in this form are required; those that are optional are listed here. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Device name	Displays the block device on which to make the file system, which corresponds to the name of the selected volume. This field is read-only and cannot be changed.
FS Type	UFS is the only currently supported file system type.
Mount point	The desired mount point for the file system. If the specified mount point does not already exist, the Visual Administrator automatically creates it. The Visual Administrator attempts to provide a default mount point, which it obtains by scanning `/etc/fstab`.
Mount automatically	Indicates whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/fstab`). No is the default.

- File System Check Form

  Basic-Ops → UFS Operations → Check File System (fsck)

  This form checks a file system that exists on a volume but is not currently mounted. The following table describes the fields for this form.

  The fields in this form are required. Fields in this form are read/write fields, unless listed as read-only.

Field	Description
Volume	Displays the name of the volume containing the file system to be checked (with `fsck`). This field is read-only and cannot be changed.
FS type	Indicates the type of the file system to be checked.

E.8.2.2 File System Properties Form

The following discussion describes the properties form. This form reveals the properties of a particular file system:

- File System Properties Form

  Basic-Ops → UFS Operations → Display Properties

  This form provides detailed information on the attributes of a particular file system. This properties form contains a list of mounted file systems, from which you can select the file system whose properties are to be displayed. The following table describes the fields for this form.

  All fields in this form are read-only and cannot be changed.

Field	Description
Mount Point	The mount point of this file system.
Device	The block device on which this file system resides.
Block Size	The block size of this file system.
Default block size	Fundamental file system block size.
Total disk space	Number of megabytes of disk storage on this file system available on the disk.
Disk space available	Number of megabytes of disk storage on this file system that are available for use.
Capacity	Percentage of the total disk storage space still available for use. This is the free space available divided by the total disk space.
Total files	The maximum number of files allowed on this file system.
Free files available	The number of files that still can be created on this file system.
FS type	The file system type (such as `UFS`).
Max file name length	The maximum number of characters for a file name on this file system. This restriction is imposed by the file system.
FS attributes	Indicates attributes associated with this file system: `read-only` indicates a file system that cannot be written to.

Glossary

The following are LSM terms and definitions.

C

concatenated plex

A plex that uses subdisks on one or more disks to create a virtual contiguous region of storage space that is accessed linearly. If LSM reaches the end of a subdisk while writing data, it continues to write data to the next subdisk, which can physically exist on the same disk or a different disk. This layout allows you to use space on several regions of the same disk, or regions of several disks, to create a single big pool of storage.

See also *RAID 5 plex*, *striped plex*

configuration database

A small database that contains all volume, plex, subdisk, and disk media records. These databases are replicated onto some or all disks in the disk group, often with two copies on each disk. Because these databases pertain to disk groups, record associations cannot span disk groups. Thus, you cannot define a subdisk on a disk in one disk group and associate it with a volume in another disk group.

D

description set

A set of files that are saved by using the `volsave` command and can be used to restore an LSM configuration. By default, an LSM description set is saved in a time-stamped directory under the `/usr/var/lsm/db` directory.

Dirty Region Log (DRL)

See *log plex*

disk

Disks exist as two entities:

- A physical disk on which all data is ultimately stored and which exhibits all the behaviors of the underlying technology.
- An LSM representation of disks which, while mapping one-to-one with the physical disks, is just a presentation of units from which allocations of storage are made.

The difference is that a physical disk presents the image of a device with a definable geometry with a definable number of cylinders, heads, and so on, while an LSM disk is simply a unit of allocation with a name and a size.

Disks used by LSM usually contain two special regions: a private region and a public region. Typically, each region is formed from a complete partition of the disk, resulting in a sliced disk; however, the private and public regions can be allocated from the same partition, resulting in a simple disk. A disk used by LSM can also be a nopriv disk, which has only a public region and no private region. Nopriv disks are created as the result of encapsulating a disk or disk partition.

See also *disk group*, *nopriv disk*, *simple disk*, *sliced disk*, *subdisk*, *volume*

disk access record

A configuration record that defines the path to a disk. Disk access records most often include a unit number. LSM uses the disk access records stored in a system to find all disks attached to the system. Disk access records do not identify particular physical disks.

Through the use of disk IDs, LSM allows you to move disks between controllers or to different locations on a controller. When you move a disk, a different disk access record is used to access the disk, although the disk media record continues to track the actual physical disk.

On some systems, LSM builds a list of disk access records automatically, based on the list of devices attached to the system. On these systems, it is not necessary to define disk access records explicitly. On other systems, you must define disk access records with the `/sbin/voldisk define` command. Specialty disks, such as RAM disks or floppy disks, are likely to require explicit `/sbin/voldisk define` commands.

Disk access records are identified by their disk access names (also known as DA names).

See also *disk ID*, *disk media record*, *volboot file*

disk group

A group of disks that share a common configuration database. A configuration database consists of a set of records describing objects including disks, volumes, plexes, and subdisks that are associated with one particular disk group. Each disk group has an administrator-assigned name that you use to reference that disk group. Each disk group has an internally defined unique disk group ID, which differentiates two disk groups with the same administrator-assigned name.

Disk groups provide a method to partition the configuration database, so that the database size is not too large and so that database modifications do not affect too many drives. They also allow LSM to operate with groups of physical disk media that can be moved between systems.

Disks and disk groups have a circular relationship: disk groups are formed from disks, and disk group configuration databases are stored on disks. All disks in a disk group are stamped with a disk group ID, which is a unique identifier for naming disk groups. Some or all disks in a disk group also store copies of the configuration database for the disk group.

See also *disk group ID*, *root disk group (rootdg)*

disk group ID

A 64–byte universally unique identifier that is assigned to a disk group when the disk group is created with the `/sbin/voldg init` command. This identifier is in addition to the disk group name, which you assigned. The disk group ID differentiates disk groups that have the same administrator-assigned name.

disk header

A block stored in a private region of a disk that defines several properties of the disk, such as the:

- Size of the private region
- Location and size of the public region
- Unique disk ID for the disk
- Disk group ID and disk group name (if the disk is currently associated with a disk group)
- Host ID for a host that has exclusive use of the disk

disk ID

A 64–byte universally unique identifier that is assigned to a physical disk when its private region is initialized with the `/sbin/voldisk init` command. The disk ID is stored in the disk media record so that the physical disk can be related to the disk media record at system startup.

See also *disk media record*

disk media record

A reference to a physical disk or possibly a disk partition. This record can be thought of as a physical disk identifier for the disk or partition. Disk media records are configuration records that provide a name (known as the disk media name or DM name) that you use to reference a particular disk independent of its location on the system's various disk controllers. Disk media records reference particular physical disks through a disk ID, which is a unique identifier that is assigned to a disk when it is initialized for use with the LSM software.

Operations are provided to set or remove the disk ID stored in a disk media record. Such operations have the effect of removing or replacing disks, with any associated subdisks being removed or replaced along with the disk.

See also *disk access record*

H

host ID

A name, usually assigned by you, that identifies a particular host. Host IDs are used to assign ownership to particular physical disks. When a disk is part of a disk group that is in active use by a particular host, the disk is stamped with that host's host ID. If another host attempts to access the disk, it detects that the disk has a nonmatching host ID and disallows access until the host with ownership discontinues use of the disk. Use the `/sbin/voldisk clearimport` command to clear the host ID stored on a disk.

If a disk is a member of a disk group and has a host ID that matches a particular host, then that host will import the disk group as part of system startup.

hot-spare disk, hot-sparing

A hot-spare disk is an LSM disk that you designate to automatically replace a disk that fails while in use by a volume. You enable and disable the hot-sparing feature with the `volwatch` command. You can designate a hot-spare disk when you initialize it for LSM use or later, as long as the disk is not being used by other LSM objects (volumes or subdisks).

K

kernel log

A log kept in the private region on the disk that is written by the LSM kernel. The log contains records describing the state of volumes in the disk group. This log provides a mechanism for the kernel to persistently register state changes, so that the `vold` daemon can detect the state changes even in the event of a system failure.

L

log plex

A plex that keeps track of write activity for a mirrored or RAID 5 volume. The log plex for a mirrored volume is called a Dirty Region Log (DRL). A DRL maintains a bitmapped representation of the regions of the volume, and marks as dirty all regions being written to. A DRL reduces the time required to restore synchronization of data for all the plexes in the volume when the system restarts after a failure. A volume can have more than one DRL for redundancy. A DRL does not provide any benefit in the event of a disk failure affecting the volume.

The log plex for a RAID 5 volume is called a RAID 5 log plex. A RAID 5 log plex maintains a bitmapped representation of the regions of the volume and marks as dirty all regions being written to. In addition, the RAID 5 log plex stores a copy of the data and parity for a predefined number of writes. In the event of a system failure, a RAID 5 log reduces the time required to restore synchronization of data for all the plexes in the volume when the system restarts. When the system restarts after a failure, the write operations that did not complete before the failure are restarted.

M

mirror

See *plex*

mirrored volume

A volume that has more than one concatenated or striped data plex and, typically, at least one log plex.

See also *RAID 5 volume*, *simple volume*

N

nopriv disk

A disk that is configured for use by LSM and has only a public region and no private region. The public region represents the space that LSM can use to create subdisks for data storage. A nopriv disk is typically created as a result of encapsulating existing data in a disk or disk partition.

See also *simple disk*, *sliced disk*

P

plex

A copy of a volume's logical data address space; also known as a mirror. A volume can have up to 32 plexes associated with it. Each plex is, at least conceptually, a copy of the volume that is maintained consistently in the presence of volume I/O and changes to the LSM configuration. Plexes represent the primary means of configuring storage for a volume. Plexes can have a concatenated, striped, or RAID 5 organization (layout).

See also *concatenated plex*, *RAID 5 plex*, *striped plex*

plex consistency

If the plexes of a volume contain different data, then the plexes are said to be inconsistent. This is a problem only if LSM is unaware of the inconsistencies, as the volume can return differing results for consecutive reads.

Plex inconsistency is a serious compromise of data integrity. This inconsistency is caused by write operations that start around the time of a system failure, if parts of the write complete on one plex but not the other. If the plexes are not first synchronized to contain the same data, plexes are inconsistent after creation of a mirrored volume. An important role of LSM is to ensure that consistent data is returned to any application that reads a volume. This might require that plex consistency of a volume be “recovered” by copying data between plexes, so that they have the same contents. Alternatively, you can put a volume into a state such that reads from one plex are automatically written back to the other plexes, making the data consistent for that volume offset.

private region

The private region of a disk contains on-disk structures that are used by LSM for internal purposes. Each private region is typically 4096 blocks long and begins with a disk header that identifies the disk and its disk group. Private regions can also contain copies of a disk group’s configuration database and copies of the disk group’s kernel log.

See also *disk header*, *kernel log*, *public region*

public region

The public region of a disk is the space reserved for allocating subdisks. Subdisks are defined with offsets that are relative to the beginning of the public region of a disk. Only one contiguous region of a disk can form the public region for a disk.

See also *private region*

R

RAID 5 plex

A plex that places data and parity evenly across each of its associated subdisks. A plex has a characteristic number of stripe columns (represented by the number of associated subdisks) and a characteristic stripe width. The stripe width defines how much data with a particular address is allocated to one of the associated subdisks. The parity data is the result of an XOR operation on the data in each stripe unit. The parity data is written to a different column (presumed to be a different disk) for each stripe, left-shifted by one column, so that no one column contains all the parity for the volume. Therefore, if a disk in a RAID 5 plex fails, the volume is still recoverable by recreating the missing data or parity for each stripe.

See also *concatenated plex*, *striped plex*

RAID 5 volume

A volume that uses a RAID 5 plex and, typically, at least one RAID 5 log plex. A RAID 5 volume has only one RAID 5 plex.

See also *mirrored volume*, *simple volume*

read policy

A configurable policy for switching between plexes for volume reads. When a volume has more than one enabled associated plex, LSM distributes reads between the plexes to distribute the I/O load and thus increase total possible bandwidth of reads through the volume. You set the read policy. Read policy choices include:

- round-robin read policy

  For every other read operation, switches to a different plex from the one used for the previous read operation. Given three plexes, switches between each of the three plexes, in order.

- preferred plex read policy

  Specifies a particular plex that is used to satisfy read requests. In the event that a read request cannot be satisfied by the preferred plex, this policy changes to the round-robin read policy.

- select read policy

  The default policy. Adjusts to use an appropriate read policy based on the set of plexes associated with the volume. If only one enabled read-write striped plex is associated with the volume, then that plex is chosen automatically as the preferred plex; otherwise, the round-robin policy is used. If a volume has one striped plex and one concatenated plex, preferring the striped plex often yields better throughput.

root disk group (rootdg)

LSM creates and requires one special disk group called rootdg. This group is generally the default for most utilities. In addition to defining the regular disk group information, the configuration database for the root disk group contains local information that is specific to a disk group.

S

simple disk

A disk that is configured for use by LSM and has a public region and a private region that occupy the same disk partition. The public region represents the space that LSM can use to create subdisks for data storage. The private region is used by LSM to store a copy of the configuration database and kernel log for the disk group to which the disk belongs. A simple disk is created by initializing a disk partition, instead of the entire disk, for LSM use.

See also *nopriv disk*, *sliced disk*

simple volume

A volume that uses only one concatenated plex. This type of volume provides no data redundancy and no protection from system failure but does permit you to create a volume using space on multiple disks (enabling creation of storage that is not bounded by disks or disk partitions), move data to other LSM disks, and perform online volume management. Without an LSM license, you can create only simple volumes.

See also *mirrored volume*, *RAID 5 volume*

sliced disk

A disk that is configured for use by LSM and has a separate public region (typically the `g` partition of the disk) and a private region (typically the `h` partition). The public region represents the space that LSM can use to create subdisks for data storage. The private region is used by LSM to store a copy of the configuration database and kernel log for the disk group to which the disk belongs.

See also *nopriv disk*, *simple disk*

striped plex

A plex that places data evenly across each of its associated subdisks. A plex has a characteristic number of stripe columns (represented by the number of associated subdisks) and a characteristic stripe width. The stripe width defines how much data with a particular address is allocated to one of the associated subdisks. Given a stripe width of 128 blocks and two stripe columns, the first group of 128 blocks is allocated to the first subdisk, the second group of 128 blocks is allocated to the second subdisk, the third group to the first subdisk, and so on.

See also *concatenated plex*, *RAID 5 plex*

subdisk

A region of storage allocated on a disk for use by a volume. Subdisks are associated with volumes through plexes. You organize one or more subdisks to form plexes based on a plex layout: concatenated, striped, or RAID 5. Subdisks are defined relative to disk media records.

V

volboot file

The `volboot` file is a special file (usually stored in `/etc/vol/volboot`) that is used to bootstrap the root disk group and to define a system's host

ID. In addition to a host ID, the `volboot` file contains a list of disk access records. On system startup, this list of disks is scanned to find a disk that is a member of the rootdg disk group and that is stamped with this system's host ID. When such a disk is found, its configuration database is read and is used to get a more complete list of disk access records that are used as a second-stage bootstrap of the root disk group and to locate all other disk groups.

See also *disk access record*

volume

A virtual disk device that looks to applications and file systems like a physical disk partition device. Volumes present block and raw device interfaces that are compatible in their use. A volume can use mirrors, span several disk drives, and be moved to different disks. You can change the configuration of a volume without causing disruption to applications or file systems that are using the volume.

Index

D

E

F

G

H

I

T

U

V

W