Sun™ StorEdge™ Volume Manager 2.6 User's Guide

THE NETWORK IS THE COMPUTER™

A Sun Microsystems, Inc. Business
901 San Antonio Road
Palo Alto, CA 94303-4900 USA
1 650 960-1300 fax 1 650 969-9131

Part No. 805-5705-10
Revision A, July 1998

Send comments about this document to: `smcc-docs@sun.com`

Copyright 1998 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94043 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, AnswerBook, SunDocs, StorEdge Volume Manager, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 1998 Sun Microsystems, Inc., 901 San Antonio Road, Palo Alto, Californie 94043 Etatis-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées des systèmes Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, AnswerBook, SunDocs, StorEdge Volume Manager, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REPONDRE A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.

Please
Recycle

Adobe PostScript

Contents

Preface

Sun StorEdge Volume Manager 2.6 User's Guide provides user information for the Sun StorEdge Volume Manager 2.6 software. These instructions are designed for an experienced system administrator.

Using UNIX Commands

This document does not contain information on basic UNIX® commands and procedures such as shutting down the system, booting the system, and configuring devices.

See one or more of the following for this information:

- *Solaris 2.x Handbook for SMCC Peripherals*
- AnswerBook™ online documentation for the Solaris™ 2.x software environment
- Other software documentation that you received with your system

Typographic Conventions

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Examples
`AaBbCc123`	The names of commands, files, and directories; on-screen computer output.	Edit your `.login` file. Use `ls -a` to list all files. `% You have mail`.
`AaBbCc123`	What you type, when contrasted with on-screen computer output.	`%` **`su`** `Password:`
AaBbCc123	Book titles, new words or terms, words to be emphasized. Command-line variable; replace with a real name or value.	Read Chapter 6 in the *User's Guide*. These are called *class* options. You *must* be `root` to do this. To delete a file, type `rm` *filename*.

Related Documentation

TABLE P-2 Related Documentation

Application	Title	Part Number
System Administration	*Sun StorEdge Volume Manager 2.6 System Administrator's Guide*	805-5706-xx
Installation	*Sun StorEdge Volume Manager 2.6 Installation Guide*	805-5707-xx
Release notes	*Sun StorEdge Volume Manager 2.6 Release Notes*	805-5708-xx
Storage Administrator GUI	*Sun StorEdge Volume Manager Storage Administrator 1.0 User's Guide*	805-5709-xx

Ordering Sun Documents

SunDocs[SM] is a distribution program for Sun Microsystems technical documentation. Contact SunExpress for easy ordering and quick delivery. You can find a listing of available Sun documentation on the World Wide Web.

TABLE P-3 SunExpress Contact Information

Country	Telephone	Fax
Belgium	02-720-09-09	02-725-88-50
Canada	1-800-873-7869	1-800-944-0661
France	0800-90-61-57	0800-90-61-58
Germany	01-30-81-61-91	01-30-81-61-92
Holland	06-022-34-45	06-022-34-46
Japan	0120-33-9096	0120-33-9097
Luxembourg	32-2-720-09-09	32-2-725-88-50
Sweden	020-79-57-26	020-79-57-27
Switzerland	0800-55-19-26	0800-55-19-27
United Kingdom	0800-89-88-88	0800-89-88-87
United States	1-800-873-7869	1-800-944-0661
World Wide Web: `http://www.sun.com/sunexpress/`		

Sun Welcomes Your Comments

Please use the *Reader Comment Card* that accompanies this document. We are interested in improving our documentation and welcome your comments and suggestions.

If a card is not available, you can email or fax your comments to us. Please include the part number of your document in the subject line of your email or fax message.

- Email: `smcc-docs@sun.com`
- Fax: SMCC Document Feedback
1-415-786-6443

CHAPTER 1

Introduction to the Volume Manager

This chapter provides detailed information about the Sun StorEdge Volume Manager. The first part of this chapter describes the Volume Manager and its features; the second part provides general information on disk arrays and Redundant Arrays of Inexpensive Disks (RAID).

1.1 Volume Manager Overview

The section describes how the Volume Manager works and the objects that VxVM manipulates. Various features of the Volume Manager are discussed later in this section.

1.1.1 How the Volume Manager Works

The Volume Manager builds virtual devices called *volumes* on top of physical disks. Volumes are accessed by a UNIX file system, a database, or other applications in the same way physical disk partitions would be accessed. Volumes are composed of other virtual objects that can be manipulated to change the volume's configuration. Volumes and their virtual components are referred to as *Volume Manager objects*. Volume Manager objects can be manipulated in a variety of ways to optimize performance, provide redundancy of data, and perform backups or other administrative tasks on one or more physical disk without interrupting applications. As a result, data availability and disk subsystem throughput are improved.

To understand the Volume Manager, you must first understand the relationships between physical objects and Volume Manager objects.

1.1.2 Physical Objects

To perform disk management tasks using the Volume Manager, you must understand two physical objects:

- Physical disks
- Partitions

1.1.2.1 Physical Disks

A *physical disk* is the underlying storage device (media), which may or may not be under Volume Manager control. A physical disk can be accessed using a device name such as `c#t#d#`, where `c#` is the controller, `t#` is the target ID, and `d#` is the disk number. The disk in FIGURE 1-1 is disk number 0 with a target ID of 0, and it is connected to controller number 0 in the system.

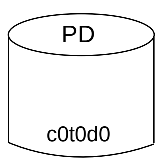

FIGURE 1-1 Example of a Physical Disk

1.1.2.2 Partitions

A physical disk can be divided into one or more *partitions*. The partition number, or `s#`, is given at the end of the device name. Note that a partition can take up an entire physical disk, such as the partition shown in FIGURE 1-2.

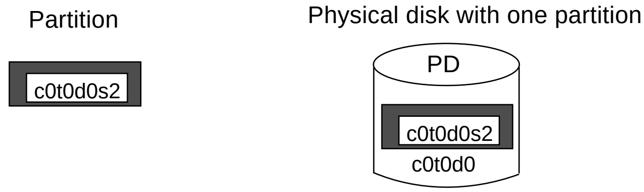

FIGURE 1-2 Example of a Partition

The relationship between physical objects and Volume Manager objects is established when you place a partition from a physical disk under Volume Manager control.

1.1.3 Volume Manager Objects

There are several Volume Manager objects that must be understood before you can use the Volume Manager to perform disk management tasks:

- VM disks
- Disk groups
- Subdisks
- Plexes
- Volumes

1.1.3.1 VM Disks

A *VM disk* is a contiguous area of disk space from which the Volume Manager allocates storage. When you place a partition from a physical disk under Volume Manager control, a VM disk is assigned to the partition. Each VM disk corresponds to at least one partition. A VM disk is typically composed of a *public region* (from which storage is allocated) and a *private region* (where configuration information is stored).

A VM disk is accessed using a unique *disk media name*, which you can supply (or else the Volume Manager assigns one that typically takes the form `disk##`). FIGURE 1-3 shows a VM disk with a disk media name of `disk01` that is assigned to the partition `c0t0d0s2`.

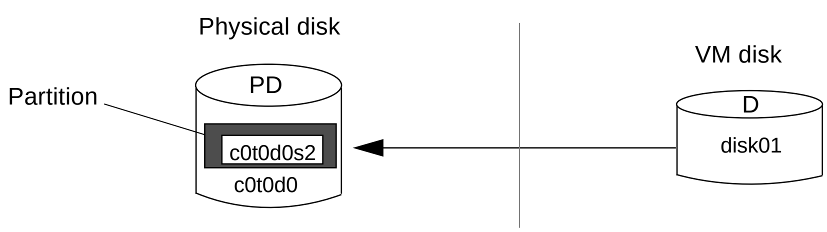

FIGURE 1-3 Example of a VM Disk

With the Volume Manager, applications access volumes (created on VM disks) rather than partitions.

1.1.3.2 Disk Groups

A *disk group* is a collection of VM disks that share a common configuration. A configuration consists of a set of records containing detailed information about existing Volume Manager objects, their attributes, and their relationships. The default disk group is `rootdg` (the root disk group). Additional disk groups can be created, as necessary. Volumes are created within a disk group; a given volume must be configured from disks belonging to the same disk group. Disk groups allow the administrator to group disks into logical collections for administrative convenience. A disk group and its components can be moved as a unit from one host machine to another.

1.1.3.3 Subdisks

A *subdisk* is a set of contiguous disk blocks; subdisks are the basic units in which the Volume Manager allocates disk space. A VM disk can be divided into one or more subdisks. Each subdisk represents a specific portion of a VM disk, which is mapped to a specific region of a physical disk. Since the default name for a VM disk is `disk`*##* (such as `disk01`), the default name for a subdisk is `disk`*##-##*. So, for example, `disk01-01` would be the name of the first subdisk on the VM disk named `disk01`.

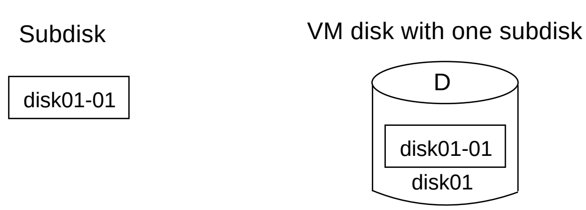

FIGURE 1-4 Example of a Subdisk

A VM disk may contain multiple subdisks, but subdisks cannot overlap or share the same portions of a VM disk. The example given in FIGURE 1-5 shows a VM disk, with three subdisks, that is assigned to one partition.

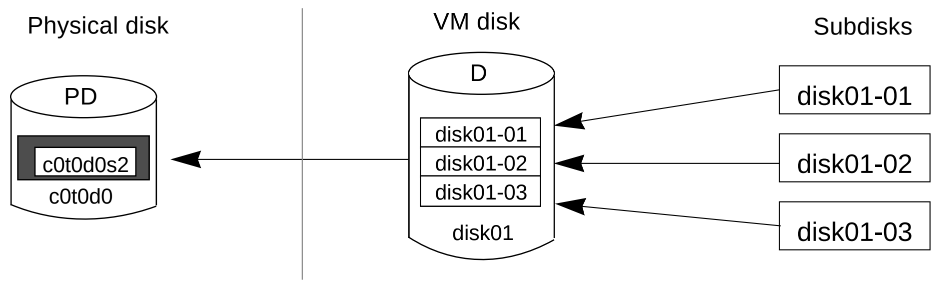

FIGURE 1-5 Example of Three Subdisks Assigned to One Partition

Any VM disk space that is not part of a subdisk is considered to be free space, which can be used to create new subdisks.

1.1.3.4 Plexes

The Volume Manager uses subdisks to build virtual entities called *plexes* (also referred to as *mirrors*). A plex consists of one or more subdisks located on one or more disks. There are three ways that data can be organized on the subdisks that constitute a plex:

- Concatenation
- Striping (RAID-0)
- RAID-5

Concatenation is discussed in this section. Details on striping (RAID-0) and RAID-5 are presented later in the chapter.

Concatenation

Concatenation maps data in a linear manner onto one or more subdisks in a plex. If you were to access all the data in a concatenated plex sequentially, you would first access the data in the first subdisk from beginning to end, then access the data in the second subdisk from beginning to end, and so forth until the end of the last subdisk.

The subdisks in a concatenated plex do not necessarily have to be physically contiguous and can belong to more than one VM disk. Concatenation using subdisks that reside on more than one VM disk is also called *spanning*.

FIGURE 1-6 illustrates concatenation with one subdisk.

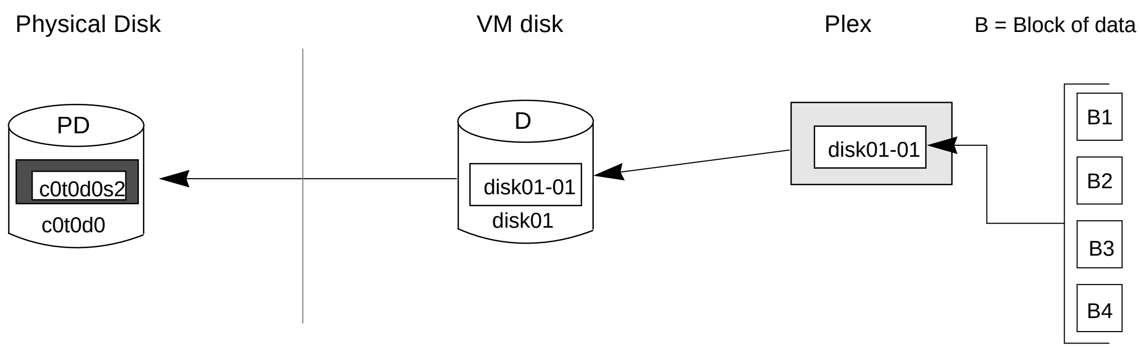

FIGURE 1-6 Example of Concatenation

Concatenation with multiple subdisks is useful when there is insufficient contiguous space for the plex on any one disk. Such concatenation can also be useful for load balancing between disks, and for head movement optimization on a particular disk.

FIGURE 1-7 shows how data would be spread over two subdisks in a spanned plex.

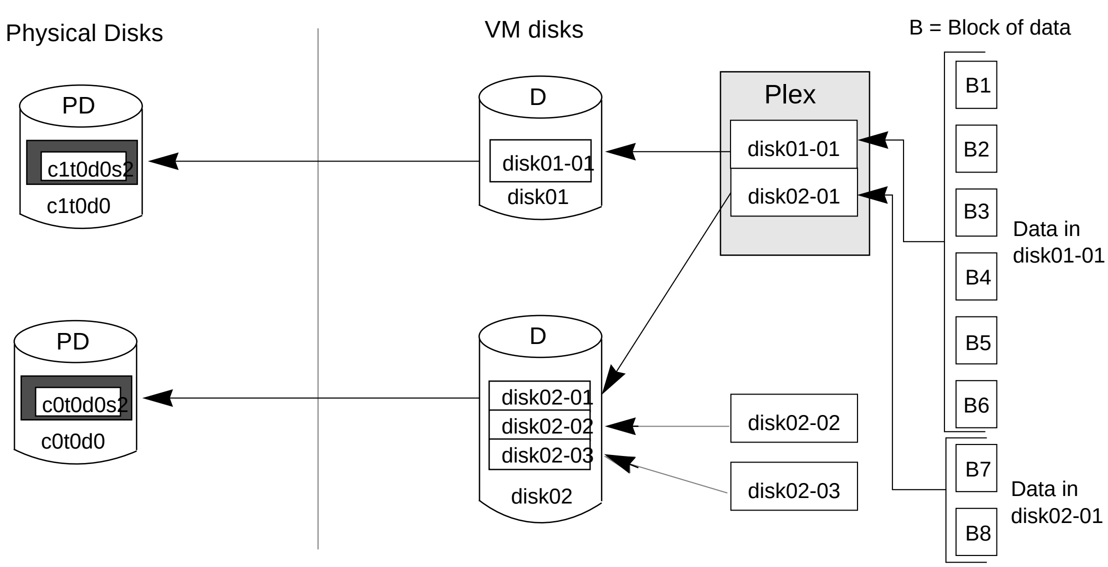

FIGURE 1-7 Example of Spanning

Since the first six blocks of data (B1 through B6) consumed most or all of the room on the partition to which VM disk disk01 is assigned, subdisk disk01-01 is alone on VM disk disk01. However, the last two blocks of data, B7 and B8, take up only a portion of the room on the partition to which VM disk disk02 is assigned. That means that the remaining free space on VM disk disk02 can be put to other uses. In this example, subdisks disk02-02 and disk02-03 are currently available for some other disk management tasks.

Caution – Spanning a plex across multiple disks increases the chance that a disk failure will result in failure of its volume. Use mirroring or RAID-5 (both described later) to substantially reduce the chance that a single disk failure will result in volume failure.

1.1.3.5 Volumes

A *volume* is a virtual disk device that appears to applications, databases, and file systems like a physical disk partition, but does not have the physical limitations of a physical disk partition. A volume consists of one or more plexes, each holding a copy of the data in the volume. Due to its virtual nature, a volume is not restricted to a particular disk or a specific area thereof. The configuration of a volume can be changed, using the Volume Manager interfaces, without causing disruption to applications or file systems that are using the volume. For example, a volume can be mirrored on separate disks or moved to use different disk storage.

A volume can consist of up to 32 plexes, each of which contains one or more subdisks. In order for a volume to be usable, it must have at least one associated plex with at least one associated subdisk. Note that all subdisks within a volume must belong to the same disk group.

The Volume Manager uses the default naming conventions of vol*##* for volumes and vol*##-##* for plexes in a volume. Administrators are encouraged to select more meaningful names for their volumes.

A volume with one plex (FIGURE 1-8) contains a single copy of the data.

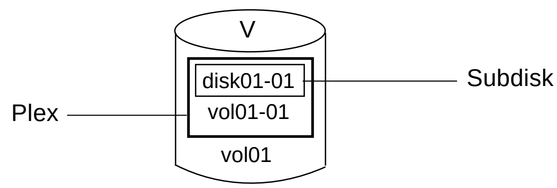

FIGURE 1-8 Example of a Volume with One Plex

Note that volume `vol01` in FIGURE 1-8 has the following characteristics:

- It contains one plex named `vol01-01`
- The plex contains one subdisk named `disk01-01`
- The subdisk `disk01-01` is allocated from VM disk `disk01`

A volume with two or more plexes (see FIGURE 1-9) is considered "mirrored" and contains mirror images of the data. See Section 1.1.5.2, "Mirroring (RAID-1)" for more information on mirrored volumes.

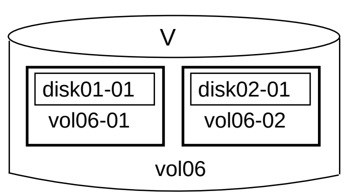

FIGURE 1-9 Example of a Volume with Two Plexes

Note that volume `vol06` in FIGURE 1-9 has the following characteristics:

- It contains two plexes named `vol06-01` and `vol06-02`
- Each plex contains one subdisk
- Each subdisk is allocated from a different VM disk (`disk01` and `disk02`)

FIGURE 1-10 shows how a volume would look if it were set up with a simple, concatenated configuration.

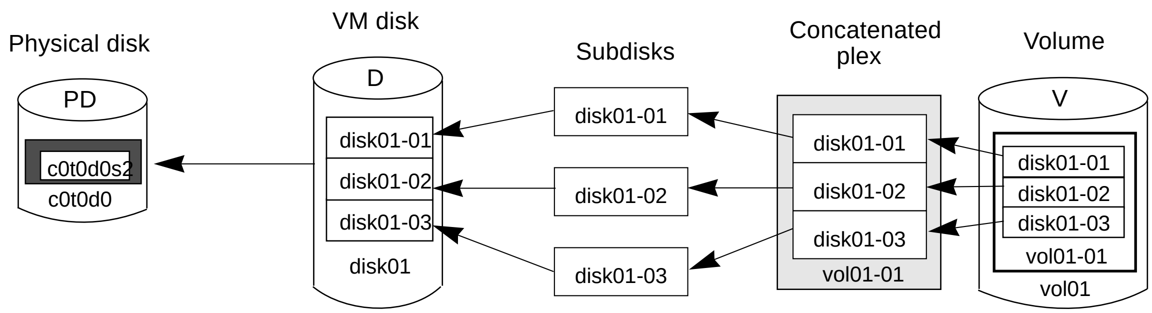

FIGURE 1-10 Example of a Volume in a Concatenated Configuration

1.1.4 Relationships Between VxVM Objects

Volume Manager objects are of little use until they are combined to build volumes. Volume Manager objects generally have the following relationship:

- VM disks are placed under VxVM control and grouped into disk groups
- One or more subdisks (each representing a specific portion of a disk) are combined to form plexes
- A volume is composed of one or more plexes

The example in FIGURE 1-11 illustrates the relationships between (virtual) Volume Manager objects, as well as how they relate to physical disks. This illustration shows a disk group containing two VM disks (`disk01` and `disk02`). `disk01` has a volume with one plex and two subdisks and `disk02` has a volume with one plex and a single subdisk.

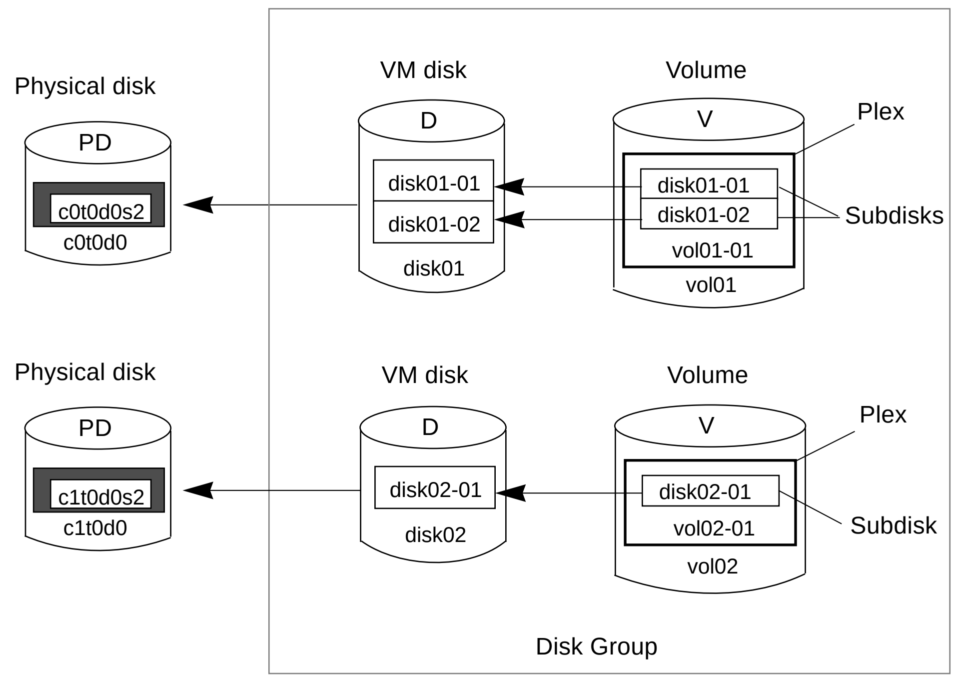

FIGURE 1-11 Relationships Between VxVM Objects

1.1.5 Volume Manager RAID Implementations

A *Redundant Array of Inexpensive Disks* (RAID) is a disk array (a group of disks that appear to the system as virtual disks, or volumes) that uses part of its combined storage capacity to store duplicate information about the data stored in the array. This duplicate information makes it possible to regenerate the data in the event of a disk failure.

This section focuses on the Volume Manager's implementations of RAID. For a general description of disk arrays and the various levels of RAID, refer to Section 1.2, "Disk Array Overview."

The Volume Manager supports the following levels of RAID:

- RAID-0 (striping)
- RAID-1 (mirroring)
- RAID-0 plus RAID-1 (striping and mirroring)
- RAID-5

The sections that follow describe how the Volume Manager implements each of these RAID levels.

1.1.5.1 Striping (RAID-0)

Striping is a technique of mapping data so that the data is interleaved among two or more physical disks. More specifically, a striped plex contains two or more subdisks, spread out over two or more physical disks. Data is allocated alternately and evenly to the subdisks of a striped plex.

The subdisks are grouped into “columns,” with each physical disk limited to one column. Each column contains one or more subdisks and can be derived from one or more physical disks. The number and sizes of subdisks per column can vary. Additional subdisks can be added to columns, as necessary.

Data is allocated in equal-sized units (called *stripe units*) that are interleaved between the columns. Each stripe unit is a set of contiguous blocks on a disk. The default stripe unit size is 64 kilobytes.

For example, if there are three columns in a striped plex and six stripe units, data is striped over three physical disks, as illustrated in FIGURE 1-12. The first and fourth stripe units are allocated in column 1; the second and fifth stripe units are allocated in column 2; and the third and sixth stripe units are allocated in column 3. Viewed in sequence, the first stripe begins with stripe unit 1 in column 1, stripe unit 2 in column 2, and stripe unit 3 in column 3. The second stripe begins with stripe unit 4 in column 1, stripe unit 5 in column 2, and stripe unit 6 in column 3. Striping continues for the length of the columns (if all columns are the same length) or until the end of the shortest column is reached. Any space remaining at the end of subdisks in longer columns becomes unused space.

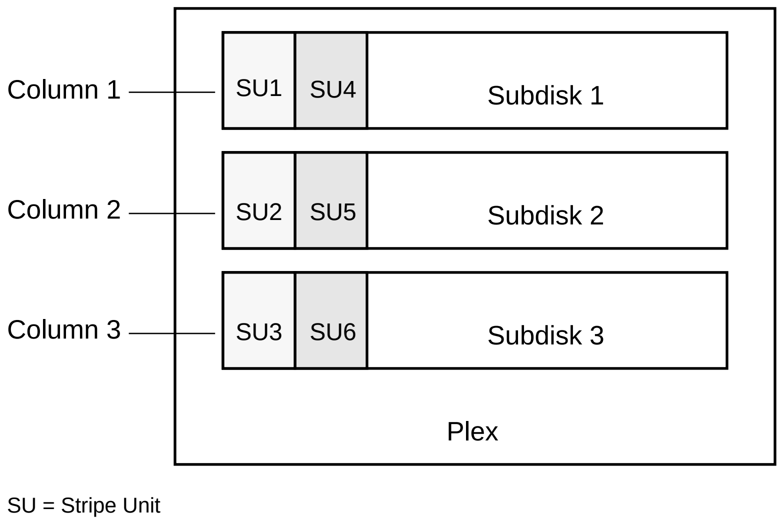

FIGURE 1-12 Striping Across Three Disks (Columns)

A *stripe* consists of the set of stripe units at the same positions across all columns. In FIGURE 1-12, stripe units 1, 2, and 3 constitute a single stripe.

Striping is useful if you need large amounts of data to be written to or read from the physical disks quickly by using parallel data transfer to multiple disks. Striping is also helpful in balancing the I/O load from multi-user applications across multiple disks.

Caution – Striping a volume, or splitting a volume across multiple disks, increases the chance that a disk failure will result in failure of that volume. For example, if five volumes are striped across the same five disks, then failure of any one of the five disks will require that all five volumes be restored from a backup. If each volume were on a separate disk, only one volume would have to be restored. Use mirroring or RAID-5 (both described later) to substantially reduce the chance that a single disk failure will result in failure of a large number of volumes.

FIGURE 1-13 shows a striped plex with three equal sized, single-subdisk columns. There is one column per physical disk.

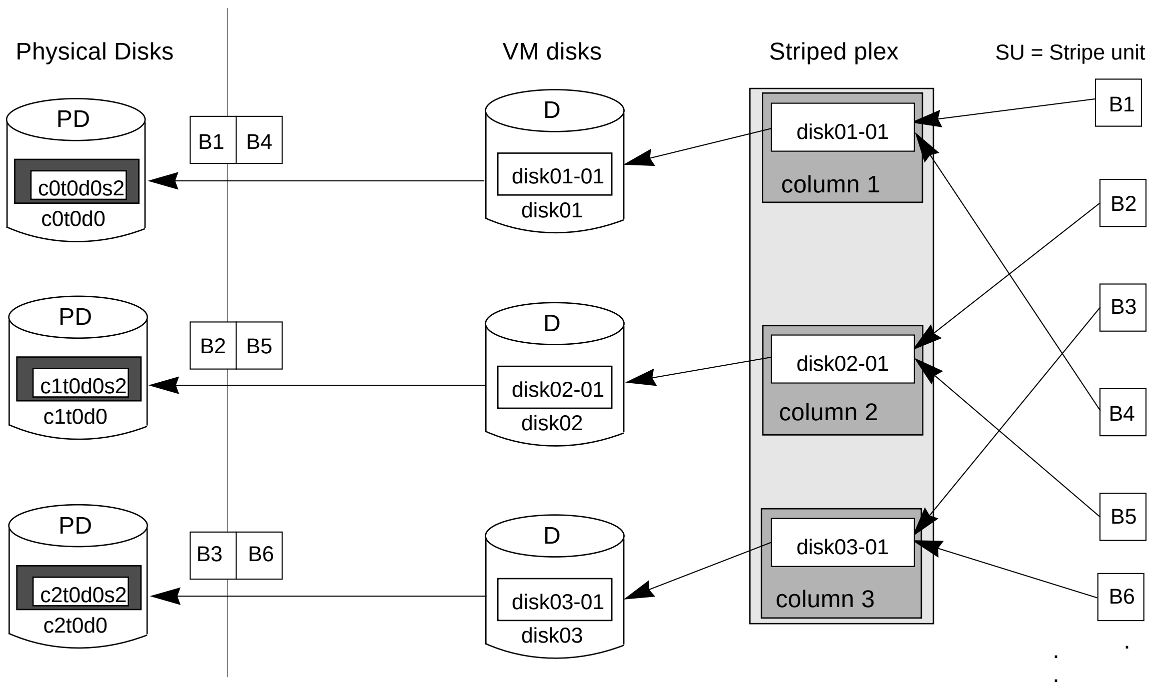

FIGURE 1-13 Example of a Striped Plex With One Subdisk per Column

Although the example in FIGURE 1-13 shows three subdisks that consume all of the VM disks, it is also possible for each subdisk in a striped plex to take up only a portion of the VM disk, thereby leaving free space for other disk management tasks.

FIGURE 1-14 shows a striped plex with 3 columns containing subdisks of different sizes. Each column contains a different number of subdisks. There is one column per physical disk. Although striped plexes are usually created using a single subdisk from each of the VM disks being striped across, it is also possible to allocate space from different regions of the same disk or from another disk (if the plex is grown, for instance).

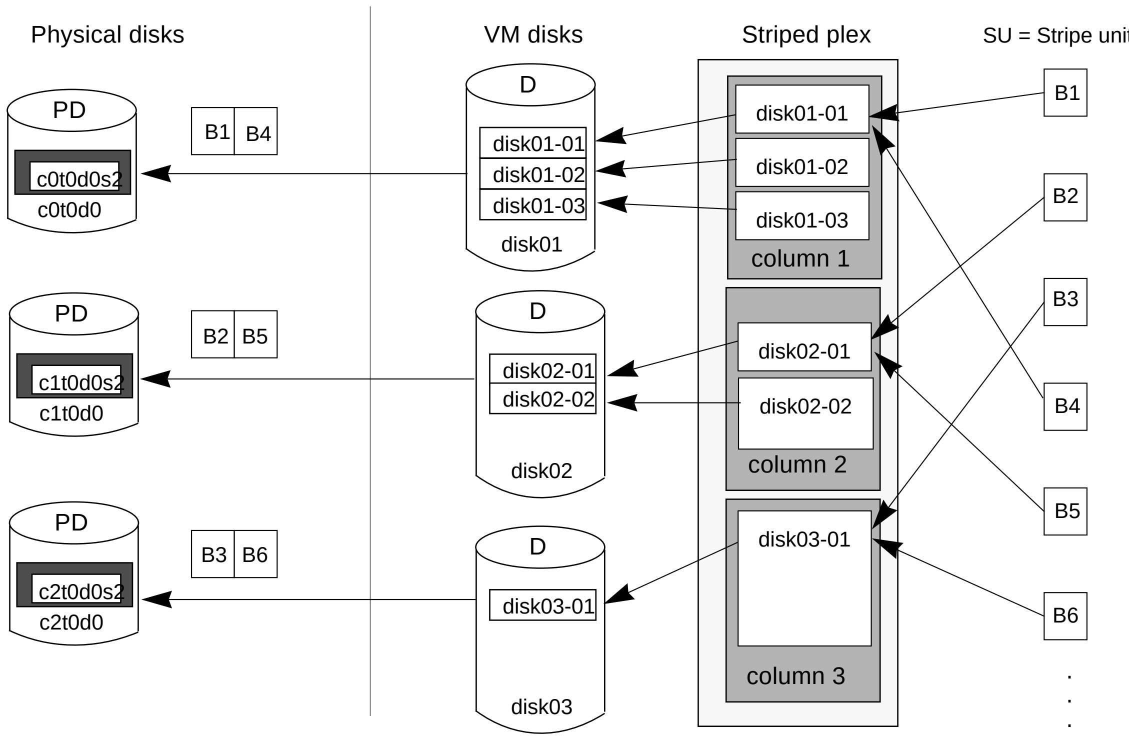

FIGURE 1-14 Example of a Striped Plex With Multiple Subdisks per Column

FIGURE 1-15 shows how a volume would look if it were set up for the simple striped configuration given in FIGURE 1-13.

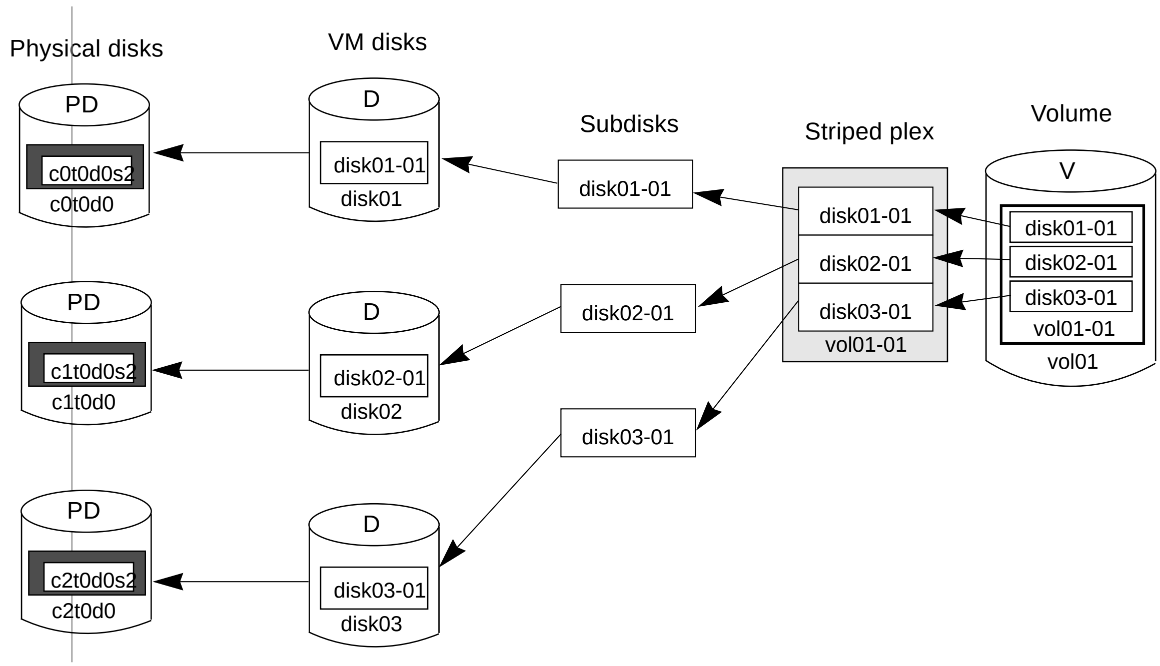

FIGURE 1-15 Example of a Volume in a Striped Configuration

1.1.5.2 Mirroring (RAID-1)

Mirroring is a technique of using multiple mirrors (plexes) to duplicate the information contained in a volume. In the event of a physical disk failure, the mirror on the failed disk becomes unavailable, but the system continues to operate using the unaffected mirrors. Although a volume can have a single plex, at least two plexes are required to provide redundancy of data. Each of these plexes should contain disk space from different disks in order for the redundancy to be effective.

When striping or spanning across a large number of disks, failure of any one of those disks will generally make the entire plex unusable. The chance of one out of several disks failing is sufficient to make it worthwhile to consider mirroring in order to improve the reliability (and availability) of a striped or spanned volume.

1.1.5.3 Striping Plus Mirroring (RAID-0 + RAID-1)

The Volume Manager supports the combination of striping with mirroring. When used together on the same volume, striping plus mirroring offers the benefits of spreading data across multiple disks while providing redundancy of data.

For striping and mirroring to be effective together, the striped plex and its mirror must be allocated from separate disks. The layout type of the mirror can be concatenated or striped.

1.1.5.4 Volume Manager and RAID-5

This section describes the Volume Manager's implementation of RAID-5. For general information on RAID-5, refer to Section 1.2.1.6, "RAID-5."

Although both mirroring (RAID-1) and RAID-5 provide redundancy of data, their approaches differ. Mirroring provides data redundancy by maintaining multiple complete copies of a volume's data. Data being written to a mirrored volume is reflected in all copies. If a portion of a mirrored volume fails, the system will continue to utilize the other copies of the data.

RAID-5 provides data redundancy through the use of *parity* (a calculated value that can be used to reconstruct data after a failure). While data is being written to a RAID-5 volume, parity is also calculated by performing an *exclusive OR* (XOR) procedure on data. The resulting parity is then written to the volume. If a portion of a RAID-5 volume fails, the data that was on that portion of the failed volume can be re-created from the remaining data and the parity.

Traditional RAID-5 Arrays

A *traditional* RAID-5 array is made up of several disks organized in rows and columns, where a *column* is a number of disks located in the same ordinal position in the array and a *row* is the minimal number of disks necessary to support the full width of a parity stripe. FIGURE 1-16 shows the row and column arrangement of a traditional RAID-5 array.

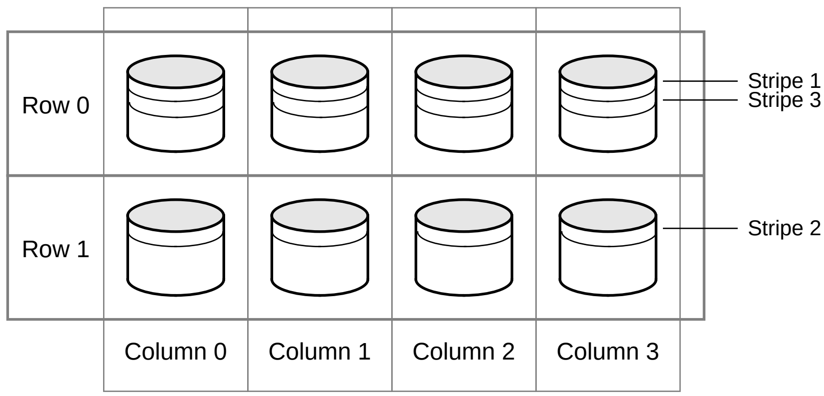

FIGURE 1-16 Traditional RAID-5 Array

This traditional array structure was developed to support growth by adding more rows per column. Striping is accomplished by applying the first stripe across the disks in Row 0, then the second stripe across the disks in Row 1, then the third stripe across Row 0's disks, and so on. This type of array requires all disks (partitions), columns, and rows to be of equal size.

VxVM RAID-5 Arrays

The Volume Manager's RAID-5 array structure differs from the traditional structure. Due to the virtual nature of its disks and other objects, the Volume Manager does not need to use rows. Instead, the Volume Manager uses columns consisting of variable length subdisks (as illustrated in FIGURE 1-17). Each subdisk represents a specific area of a disk.

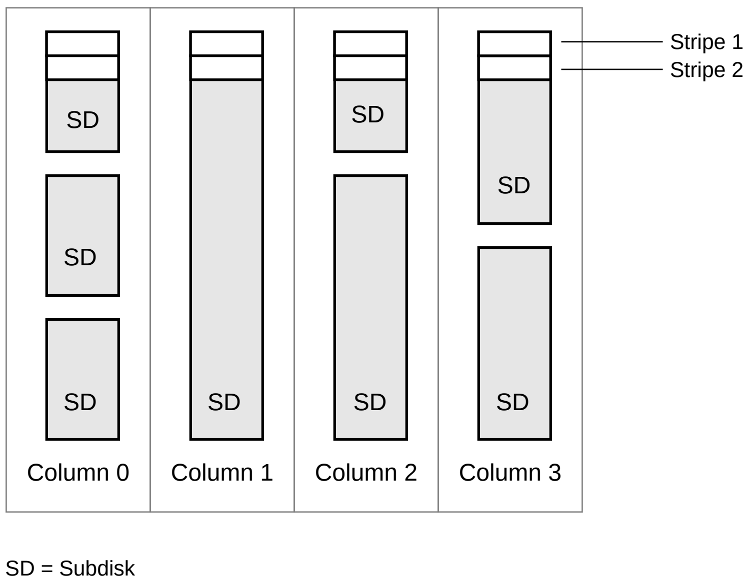

FIGURE 1-17 VxVM RAID-5 Array

With the Volume Manager RAID-5 array structure, each column can consist of a different number of subdisks and the subdisks in a given column can be derived from different physical disks. Additional subdisks can be added to the columns, as necessary. Striping (described in Section 1.1.5.1, “Striping (RAID-0)”) is accomplished by applying the first stripe across each subdisk at the top of each column, then another stripe below that, and so on for the entire length of the columns. For each stripe, an equal-sized stripe unit is placed in each column. With RAID-5, the default stripe unit size is 16 kilobytes.

Note – Mirroring of RAID-5 volumes is not currently supported.

Left-Symmetric Layout

There are several layouts for data and parity that can be used in the setup of a RAID-5 array. The layout selected for the Volume Manager’s implementation of RAID-5 is the left-symmetric layout. The left-symmetric parity layout provides optimal performance for both random I/Os and large sequential I/Os. In terms of performance, the layout selection is not as critical as the number of columns and the stripe unit size selection.

The left-symmetric layout stripes both data and parity across columns, placing the parity in a different column for every stripe of data. The first parity stripe unit is located in the rightmost column of the first stripe. Each successive parity stripe unit is located in the next stripe, left-shifted one column from the previous parity stripe unit location. If there are more stripes than columns, the parity stripe unit placement begins in the rightmost column again.

FIGURE 1-18 illustrates a left-symmetric parity layout consisting of five disks (one per column).

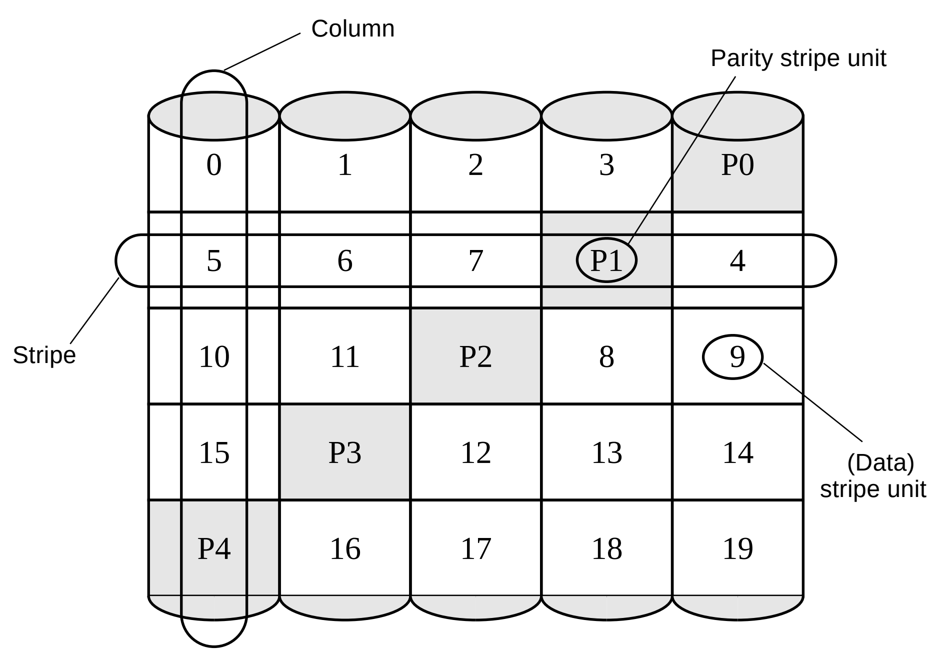

FIGURE 1-18 Left-Symmetric Layout

For each stripe, data is organized starting to the right of the parity stripe unit. In FIGURE 1-18, data organization for the first stripe begins at P0 and continues to stripe units 0-3. Data organization for the second stripe begins at P1, then continues to stripe unit 4, and on to stripe units 5-7. Data organization proceeds in this manner for the remaining stripes.

Each parity stripe unit contains the result of an exclusive OR (XOR) procedure performed on the data in the data stripe units within the same stripe. If data on a disk corresponding to one column is inaccessible due to hardware or software failure, data can be restored by XORing the contents of the remaining columns' data stripe units against their respective parity stripe units (for each stripe). For example, if the disk corresponding to the leftmost column in FIGURE 1-18 were to fail, the volume would be placed in a degraded mode. While in degraded mode, the data from the failed column could be re-created by XORing stripe units 1-3 against parity stripe unit P0 to re-create stripe unit 0, then XORing stripe units 4, 6, and 7 against parity stripe unit P1 to recreate stripe unit 5, and so on.

Note – Failure of multiple columns in a plex with a RAID-5 layout will detach the volume. This means that the volume will no longer be allowed to satisfy read or write requests. Once the failed columns have been recovered, it might be necessary to recover the user data from backups.

Logging

Without logging, it is possible for data not involved in any active writes to be lost or silently corrupted if a disk fails and the system also fails. If this double-failure occurs, there is no way of knowing if the data being written to the data portions of the disks or the parity being written to the parity portions have actually been written. Therefore, the recovery of the corrupted disk may be corrupted itself.

Logging is used to prevent this corruption of recovery data. A log of the new data and parity is made on a persistent device (such as a disk-resident volume or non-volatile RAM). The new data and parity are then written to the disks.

In FIGURE 1-19, the recovery of Disk B is dependent on the data on Disk A and the parity on Disk C having both completed. The diagram shows a completed data write and an incomplete parity write causing an incorrect data reconstruction for the data on Disk B.

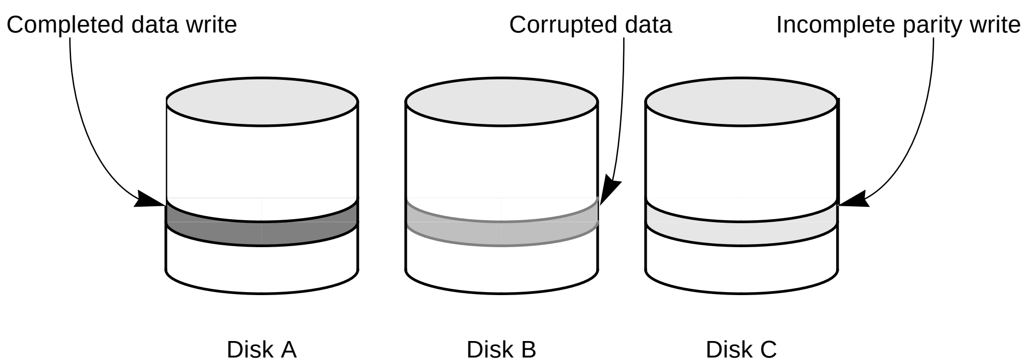

FIGURE 1-19 Incomplete Write

This failure case can be handled by logging all data writes before committing them to the array. In this way, the log can be replayed, causing the data and parity updates to be completed before the reconstruction of the failed drive takes place.

Logs are associated with a RAID-5 volume by being attached as additional, non-RAID-5 layout plexes. More than one log plex can exist per RAID-5 volume, in which case the log areas are mirrored.

Read-Modify-Write

When you write to a RAID-5 array, the following steps may be followed for each stripe involved in the I/O:

1. The data stripe units to be updated with new write data are accessed and read into internal buffers. The parity stripe unit is read into internal buffers.
2. The parity is updated to reflect the contents of the new data region. First, the contents of the old data undergo an exclusive OR (XOR) with the parity (logically removing the old data). The new data is then XORed into the parity (logically adding the new data). The new data and new parity are written to a log.
3. The new parity is written to the parity stripe unit. The new data is written to the data stripe units. All stripe units are written in a single write.

This process is known as a *read-modify-write* cycle, which is the default type of write for RAID-5. If a disk fails, both data and parity stripe units on that disk become unavailable. The disk array is then said to be operating in a *degraded* mode.

The read-modify-write sequence is illustrated in FIGURE 1-20.

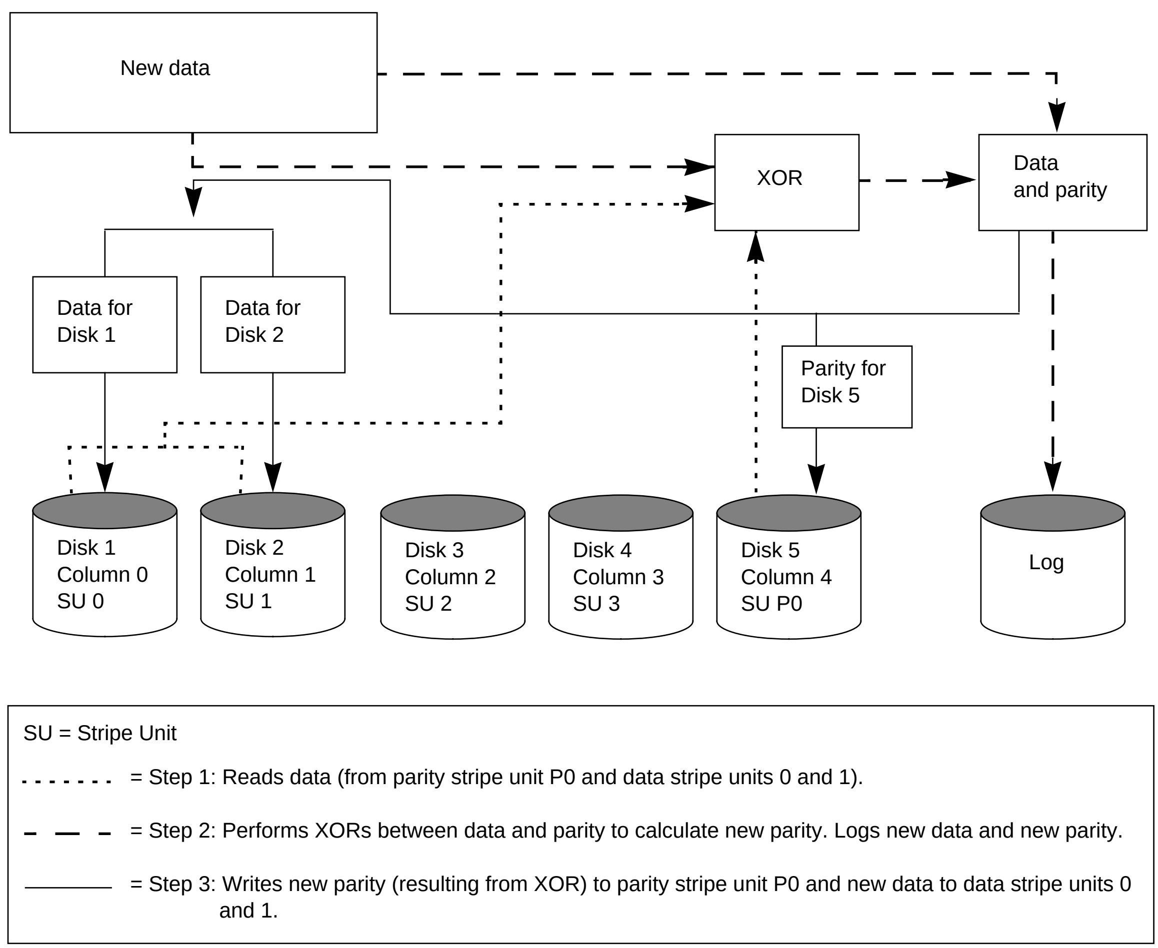

FIGURE 1-20 Read-Modify-Write

Full-Stripe Writes

When large writes (writes that cover an entire data stripe) are issued, the read-modify-write procedure can be bypassed in favor of a *full-stripe write*. A full-stripe write is faster than a read-modify-write because it does not require the read process to take place. Eliminating the read cycle reduces the I/O time necessary to write to the disk.

A full-stripe write procedure consists of the following steps:

1. All the new data stripe units are XORed together, generating a new parity value. The new data and new parity is written to a log.
2. The new parity is written to the parity stripe unit. The new data is written to the data stripe units. The entire stripe is written in a single write.

FIGURE 1-21 shows a full-stripe write.

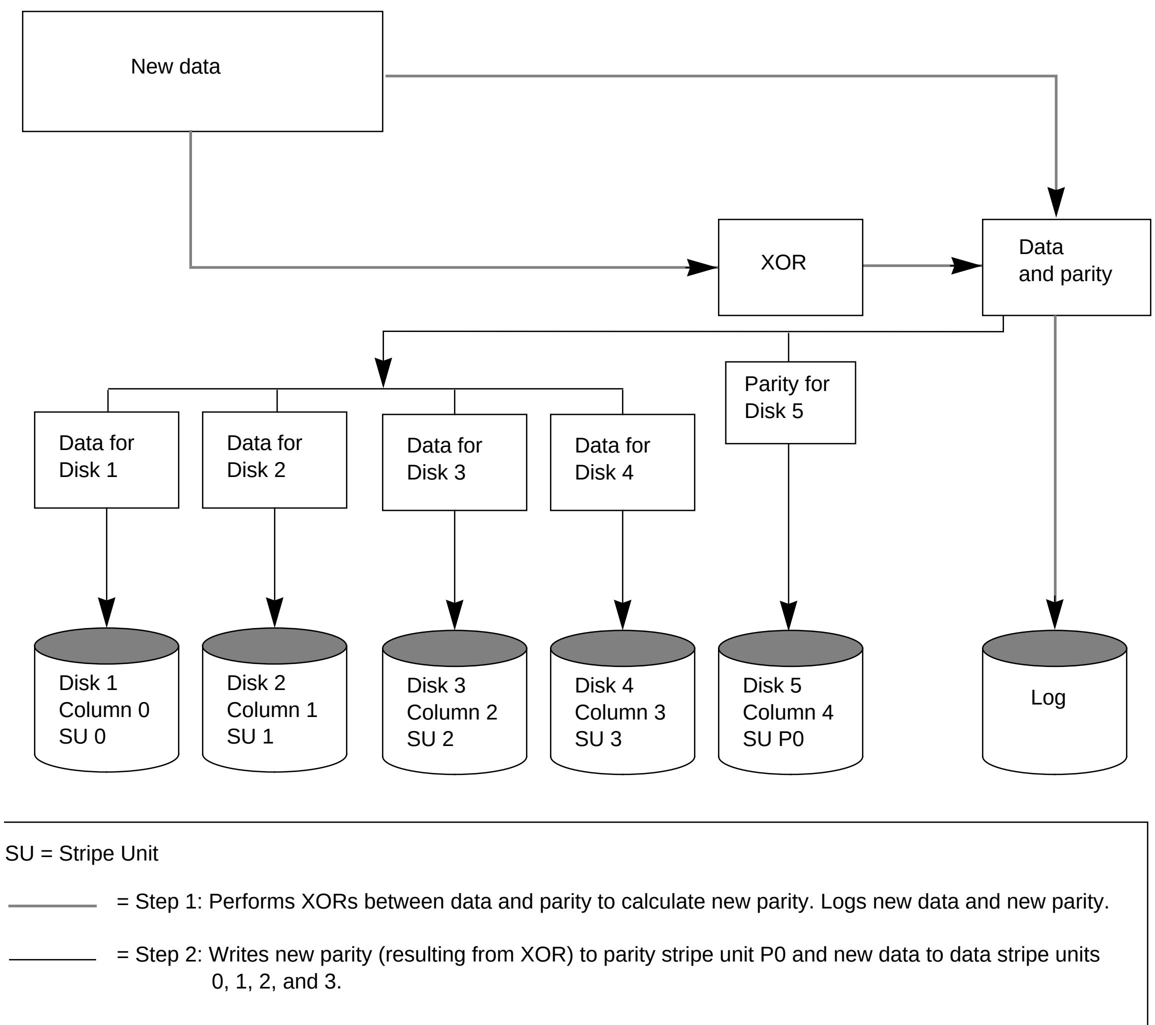

FIGURE 1-21 Full-Stripe Write

Reconstruct-Writes

When 50 percent or more of the data disks are undergoing writes in a single I/O, a *reconstruct-write* can be used. A reconstruct-write saves I/O time by XORing because it does not require a read of the parity region and only requires a read of the unaffected data (which amounts to less than 50 percent of the stripe units in the stripe).

A reconstruct-write procedure consists of the following steps:

1. Unaffected data is read from the unchanged data stripe unit(s).
2. The new data is XORed with the old, unaffected data to generate a new parity stripe unit. The new data and resulting parity are logged.
3. The new parity is written to the parity stripe unit. The new data is written to the data stripe units. All stripe units are written in a single write.

FIGURE 1-22 illustrates a reconstruct-write. A reconstruct-write is preferable to a read-modify-write in this situation because it reads only the necessary data disks, rather than reading the disks and the parity disk.

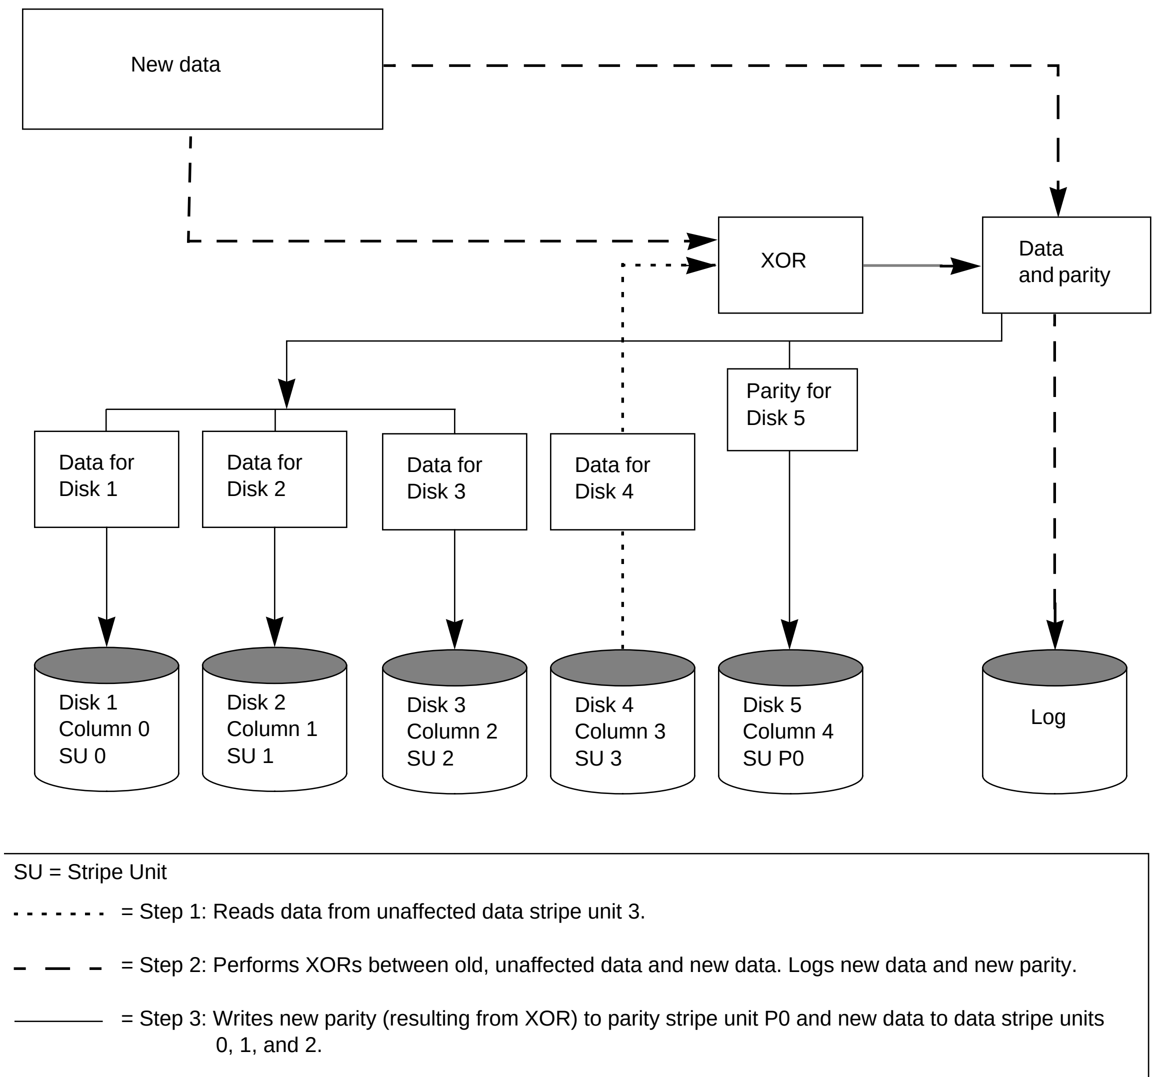

FIGURE 1-22 Reconstruct-Write

1.1.6 Hot-Relocation

Hot-relocation is the ability of a system to automatically react to I/O failures on redundant (mirrored or RAID-5) VxVM objects and restore redundancy and access to those objects. The Volume Manager detects I/O failures on VxVM objects and relocates the affected subdisks to disks designated as *spare disks* and/or free space within the disk group. The Volume Manager then reconstructs the VxVM objects that existed before the failure and makes them redundant and accessible again.

When a partial disk failure occurs (that is, a failure affecting only some subdisks on a disk), redundant data on the failed portion of the disk is relocated and the existing volumes comprised of the unaffected portions of the disk remain accessible.

Note – Hot-relocation is only performed for redundant (mirrored or RAID-5) subdisks on a failed disk. Non-redundant subdisks on a failed disk are not relocated, but the system administrator is notified of their failure.

1.1.6.1 How Hot-Relocation Works

The hot-relocation feature is enabled by default. No system administrator intervention is needed to start hot-relocation when a failure occurs.

The hot-relocation daemon, `vxrelocd`, is responsible for monitoring VxVM for events that affect redundancy and performing hot-relocation to restore redundancy. `vxrelocd` also notifies the system administrator (via electronic mail) of failures and any relocation and recovery actions. See the `vxrelocd`(1M) manual page for more information on `vxrelocd`.

The `vxrelocd` daemon starts during system startup and monitors the Volume Manager for failures involving disks, plexes, or RAID-5 subdisks. When such a failure occurs, it triggers a hot-relocation attempt.

A successful hot-relocation process involves:

1. Detecting VxVM events resulting from the failure of a disk, plex, or RAID-5 subdisk.
2. Notifying the system administrator (and other users designated for notification) of the failure and identifying the affected VxVM objects. This is done through electronic mail.
3. Determining which subdisks can be relocated, finding space for those subdisks in the disk group, and relocating the subdisks. Notifying the system administrator of these actions and their success or failure.

4. Initiating any recovery procedures necessary to restore the volumes and data. Notifying the system administrator of the recovery's outcome.

Note – Hot-relocation does not guarantee the same layout of data or the same performance after relocation. The system administrator may therefore want to make some configuration changes after hot-relocation occurs.

1.1.6.2 How Space Is Chosen for Relocation

A spare disk must be initialized and placed in a disk group as a spare *before* it can be used for replacement purposes. If no disks have been designated as spares when a failure occurs, VxVM automatically uses any available free space in the disk group in which the failure occurs. If there is not enough spare disk space, a combination of spare space and free space is used. The system administrator can designate one or more disks as hot-relocation spares within each disk group. Disks can be designated as spares using the Visual Administrator, `vxdiskadm`, or `vxedit` (as described in later chapters). Disks designated as spares do not participate in the free space model and should not have storage space allocated on them.

When selecting space for relocation, hot-relocation preserves the redundancy characteristics of the VxVM object that the relocated subdisk belongs to. For example, hot-relocation ensures that subdisks from a failed plex are not relocated to a disk containing a mirror of the failed plex. If redundancy cannot be preserved using any available spare disks and/or free space, hot-relocation does not take place. If relocation is not possible, the system administrator is notified and no further action is taken.

When hot-relocation takes place, the failed subdisk is removed from the configuration database and VxVM takes precautions to ensure that the disk space used by the failed subdisk is not recycled as free space.

For information on how to take advantage of hot-relocation, refer to Chapter 2 and Chapter 3 of the *Sun StorEdge Volume Manager System Administrator's Guide*. Refer to the *Sun StorEdge Volume Manager Installation Guide* for information on how to disable hot-relocation.

1.1.7 Volume Resynchronization

When storing data redundantly using mirrored or RAID-5 volumes, the Volume Manager takes necessary measures to ensure that all copies of the data match exactly. However, under certain conditions (usually due to complete system failures), small amounts of the redundant data on a volume can become inconsistent or *unsynchronized*. Aside from normal configuration changes (such as detaching and reattaching a plex), this can only occur when a system crashes while data is being written to a volume. Data is written to the mirrors of a volume in parallel, as is the data and parity in a RAID-5 volume. If a system crash occurs before all the individual writes complete, it is possible for some writes to complete while others do not, resulting in the data becoming unsynchronized. This is very undesirable. For mirrored volumes, it can cause two reads from the same region of the volume to return different results if different mirrors are used to satisfy the read request. In the case of RAID-5 volumes, it can lead to parity corruption and incorrect data reconstruction.

When the Volume Manager recognizes this situation, it needs to make sure that all mirrors contain exactly the same data and that the data and parity in RAID-5 volumes agree. This process is called *volume resynchronization*. For volumes that are part of disk groups that are automatically imported at boot time (such as `rootdg`), the resynchronization process takes place when the system reboots.

Not all volumes may require resynchronization after a system failure. Volumes that were never written or that were quiescent (that is., had no active I/O) when the system failure occurred could not have had any outstanding writes and thus do not require resynchronization. The Volume Manager notices when a volume is first written and marks it as *dirty*. When a volume is closed by all processes or stopped cleanly by the administrator, all writes will have completed and the Volume Manager removes the dirty flag for the volume. Only volumes that are marked dirty when the system reboots require resynchronization.

The exact process of resynchronization depends on the type of volume. RAID-5 volumes that contain RAID-5 logs can simply “replay” those logs. If no logs are available, the volume is placed in reconstruct-recovery mode and all parity is regenerated. For mirrored volumes, resynchronization is achieved by placing the volume in recovery mode (also called *read-writeback recovery mode*) and resynchronizing all data in the volume in the background. This allows the volume to be available for use while recovery is ongoing.

The process of resynchronization can be computationally expensive and can have a significant impact on system performance. The recovery process attempts to alleviate some of this impact by attempting to “spread out” recoveries to avoid stressing a specific disk or controller. Additionally, for very large volumes or for a very large number of volumes, the resynchronization process can take a long time. These effects can be addressed by using Dirty Region Logging for mirrored volumes, or by making sure that RAID-5 volumes have valid RAID-5 logs. For volumes used by database applications, the VxSmartSync Recovery Accelerator can be used (see Section 1.1.9, “VxSmartSync Recovery Accelerator”).

1.1.8 Dirty Region Logging

Dirty Region Logging (DRL) is an optional property of a volume, used to provide a speedy recovery of mirrored volumes after a system failure. DRL keeps track of the regions that have changed due to I/O writes to a mirrored volume and uses this information to recover only the portions of the volume that need to be recovered.

If DRL is not used and a system failure occurs, all mirrors of the volumes must be restored to a consistent state by copying the full contents of the volume between its mirrors. This process can be lengthy and I/O intensive; it may also be necessary to recover the areas of volumes that are already consistent.

DRL logically divides a volume into a set of consecutive regions. It keeps track of volume regions that are being written to. A dirty region log is maintained that contains a status bit representing each region of the volume. For any write operation to the volume, the regions being written are marked dirty in the log before the data is written. If a write causes a log region to become dirty when it was previously clean, the log is synchronously written to disk before the write operation can occur. On system restart, the Volume Manager will recover only those regions of the volume that are marked as dirty in the dirty region log.

Log subdisks are used to store the dirty region log of a volume that has DRL enabled. A volume with DRL has at least one log subdisk; multiple log subdisks can be used to mirror the dirty region log. Each log subdisk is associated with one of the volume's plexes. Only one log subdisk can exist per plex. If the plex contains only a log subdisk and no data subdisks, that plex can be referred to as a *log plex*. The log subdisk can also be associated with a regular plex containing data subdisks, in which case the log subdisk risks becoming unavailable in the event that the plex must be detached due to the failure of one of its data subdisks.

If the `vxassist` command is used to create a dirty region log, it creates a log plex containing a single log subdisk, by default. A dirty region log can also be created "manually" by creating a log subdisk and associating it with a plex. In the latter case, the plex may contain both a log subdisk and data subdisks.

Only a limited number of bits can be marked dirty in the log at any time. The dirty bit for a region is not cleared immediately after writing the data to the region. Instead, it remains marked as dirty until the corresponding volume region becomes the least recently used. If a bit for a given region is already marked dirty and another write to the same region occurs, it is not necessary to write the log to the disk before the write operation can occur.

Note – DRL adds a small I/O overhead for most write access patterns.

1.1.9 VxSmartSync Recovery Accelerator

The VxSmartSync Recovery Accelerator for Mirrored Oracle Databases from Veritas is a collection of features designed to speed up the resynchronization process (also known as *resilvering*) for volumes used in conjunction with the Oracle Universal Database. These features employ an extended interface between VxVM volumes and the database software so that they can cooperate to avoid unnecessary work during mirror resynchronization. These extensions can result in an order of magnitude improvement in volume recovery times.

The only requirement to take advantage of VxSmartSync is to correctly configure the volumes. Once configured, no further user action is necessary. In the Volume Manager's view, there are two types of volumes used by the database:

- *redo log volumes* contain a database's redo logs
- *data volumes* are all other volumes used by the database (control files and tablespace files)

VxSmartSync handles these two types of volumes differently, and they must be configured correctly to take full advantage of the extended interfaces. The only difference between the two types of volumes is that redo log volumes should have dirty region logs, while data volumes should not.

1.1.9.1 Data Volume Configuration

The improvements in recovery time for data volumes are achieved by letting the database software decide exactly which portions of the volume require recovery. The database keeps its own logs of changes to the data in the database; therefore, it "knows" exactly which portions of the volume require recovery. The overall recovery time is significantly reduced by reducing the amount of space that requires recovery and enabling the database to control the recovery process. Additionally, the recovery takes place when the database software is started, not at system startup; this reduces the overall impact of recovery when the system reboots. Because the recovery is controlled by the database, the recovery time for the volume is exactly the resilvering time for the database (that is., the time required to replay the redo logs).

Because the database keeps its own logs, it is not necessary for VxVM to do any logging. Data volumes should therefore be configured as mirrored volumes *without* dirty region logs. In addition to improving recovery time, this avoids any run-time I/O overhead due to DRL, which improves normal database write access.

1.1.9.2 Redo Log Volume Configuration

A *redo log* is a log of changes to the database data. No logs of the changes to the redo logs are kept by the database, so the database itself cannot provide information about which sections require resilvering. Redo logs are also written sequentially, and since traditional dirty region logs are most useful with randomly-written data, they are of minimal use for reducing recovery time for redo logs. However, VxVM can significantly reduce the number of dirty regions by modifying the behavior of its Dirty Region Logging feature to take advantage of sequential access patterns. This decreases the amount of data needing recovery and significantly reduces recovery time impact on the system.

The enhanced interfaces for redo logs enable the database software to inform VxVM when a volume is to be used as a redo log. This enables VxVM to modify the volume's DRL behavior to take advantage of the access patterns. Since the improved recovery time depends on dirty region logs, redo log volumes should be configured as mirrored volumes *with* dirty region logs.

1.1.10 Volume Manager Rootability

The Volume Manager provides the capability of placing the root file system, `swap` device, and `usr` file system under Volume Manager control — this is called *rootability*. The *root disk* (that is, the disk containing the root file system) can be put under VxVM control through the process of *encapsulation*, which converts existing partitions on that disk to volumes. Once under VxVM control, the `root` and `swap` devices appear as volumes and provide the same characteristics as other VxVM volumes. A volume that is configured for use as a swap area is referred to as a *swap volume*; a volume that contains the root file system is referred to as a *root volume*.

It is possible to mirror the `rootvol` and `swapvol` volumes, as well as other parts of the root disk required for a successful boot of the system (such as `/usr`). This provides complete redundancy and recovery capability in the event of disk failure. Without Volume Manager rootability, the loss of the `root`, `swap`, or `usr` partition would prevent the system from being booted from surviving disks.

Mirroring disk drives critical to booting ensures that no single disk failure will leave the system unusable. Therefore, a suggested configuration would be to mirror the critical disk onto another available disk (using the `vxdiskadm` command). If the disk containing the `root` and `swap` partitions fails, the system can be rebooted from the disk containing the root mirror. For more information on mirroring the boot (root) disk and system recovery procedures, see the "Recovery" appendix in the *Sun StorEdge Volume Manager System Administrator's Guide*.

1.1.10.1 Booting With Root Volumes

Ordinarily, when the operating system is booted, the root file system and `swap` area need to be available for use very early in the boot procedure (which is long before user processes can be run to load the Volume Manager configuration and start volumes). The `root` and `swap` device configurations must be completed prior to starting the Volume Manager. Starting VxVM's `vxconfigd` daemon as part of the `init` process is too late to configure volumes for use as a `root` or `swap` device.

To circumvent this restriction, the mirrors of the `rootvol` and `swapvol` volumes can be accessed by the system during startup. During startup, the system sees the `rootvol` and `swapvol` volumes as regular partitions and accesses them using standard partition numbering. Therefore, `rootvol` and `swapvol` volumes must be created from contiguous disk space that is also mapped by a single partition for each. Due to this restriction, it is not possible to stripe or span the primary plex (that is, the plex used for booting) of a `rootvol` or `swapvol` volume. Similarly, any mirrors of these volumes that might need to be used for booting cannot be striped or spanned.

1.1.10.2 Boot-Time Volume Restrictions

The `rootvol` and `usr` volumes differ from other volumes in that they have very specific restrictions on the configuration of the volumes:

- The root volume (`rootvol`) must exist in the default disk group, `rootdg`. Although other volumes named `rootvol` may be created in disk groups other than `rootdg`, only the `rootvol` in `rootdg` can be used to boot the system.
- A `rootvol` volume has a specific minor device number: minor device 0. The `usr` volume does not have a specific minor device number.
- Restricted mirrors of `rootvol`, `var`, and `usr` devices will have "overlay" partitions created for them. An "overlay" partition is one that exactly encompasses the disk space occupied by the restricted mirror. During boot (before the `rootvol`, `var`, and `usr` volumes are fully configured), the default volume configuration uses the overlay partition to access the data on the disk. (See Section 1.1.10.1, "Booting With Root Volumes.")
- Although it is possible to add a striped mirror to a `rootvol` device for performance reasons, you cannot stripe the primary plex or any mirrors of `rootvol` that may be needed for system recovery or booting purposes if the primary plex fails.
- `rootvol` and `swapvol` cannot be spanned or contain a primary plex with multiple non-contiguous subdisks.
- When mirroring parts of the boot disk, the disk being mirrored to must be large enough to hold the data on the original plex, or mirroring may not work.
- `rootvol` and `usr` cannot be DRL volumes.

In addition to these requirements, it is a good idea to have at least one contiguous mirror for each of the volumes for `root`, `usr`, `var`, `opt`, `varadm`, `usrkvm`, and `swap`. This makes it easier to convert these from volumes back to regular disk partitions (during an operating system upgrade, for example).

1.1.11 Volume Manager Daemons

Two daemons must be running in order for the Volume Manager to work properly:

- `vxconfigd`
- `vxiod`

1.1.11.1 The Volume Manager Configuration Daemon

The Volume Manager configuration daemon (`vxconfigd`) is responsible for maintaining Volume Manager disk and disk group configurations. `vxconfigd` communicates configuration changes to the kernel and modifies configuration information stored on disk. `vxconfigd` must be running before Volume Manager operations can be performed.

Starting the Volume Manager Configuration Daemon

`vxconfigd` is invoked by startup scripts during the boot procedure.

To determine whether the volume daemon is enabled, type the following command:

```
vxdctl mode
```

The following message appears if `vxconfigd` is both running and enabled:

```
mode: enabled
```

The following message appears if `vxconfigd` is running, but not enabled:

```
mode: disabled
```

To enable the volume daemon, type:

```
vxdctl enable
```

The following message appears if `vxconfigd` is not running:

```
mode: not-running
```

If the latter message appears, start `vxconfigd` by typing:

```
vxconfigd
```

Once started, `vxconfigd` automatically becomes a background process.

By default, `vxconfigd` issues errors to the console. However, `vxconfigd` can be configured to log errors to a log file.

For more information on the `vxconfigd` daemon, refer to the `vxconfigd`(1M) and `vxdctl`(1M) manual pages.

1.1.11.2 The Volume I/O Daemon

The volume extended I/O daemon (`vxiod`) allows for some extended I/O operations without blocking calling processes.

For more detailed information about `vxiod`, refer to the `vxiod` (1M) manual page.

Starting the Volume I/O Daemon

`vxiod` daemons are started at system boot time. There are typically several `vxiod` daemons running at all times. Rebooting after your initial installation should start `vxiod`.

Verify that `vxiod` daemons are running by typing:

```
vxiod
```

Since `vxiod` is a kernel thread and is not visible to users via `ps`, this is the only way to see if any `vxiod`s are running.

If any vxiod daemons are running, the following should be displayed:

```
10 volume I/O daemons running
```

where 10 is the number of vxiod daemons currently running.

If no vxiod daemons are currently running, start some by typing:

```
vxiod set 10
```

where 10 may be substituted by the desired number of vxiod daemons. It is generally recommended that at least one vxiod daemon exist per CPU.

1.1.12 Volume Manager Interfaces

The Volume Manager supports the following user interfaces:

- Visual Administrator — The Visual Administrator is a graphical user interface to the Volume Manager. The Visual Administrator provides visual elements such as icons, menus, and forms to ease the task of manipulating Volume Manager objects. In addition, the Visual Administrator acts as an interface to some common file system operations.
- Command Line Interface — The Volume Manager command set consists of a number of comprehensive commands that range from simple commands requiring minimal user input to complex commands requiring detailed user input. Many of the Volume Manager commands require a thorough understanding of Volume Manager concepts. Most Volume Manager commands require superuser privileges.
- Volume Manager Support Operations — The Volume Manager Support Operations interface (vxdiskadm) provides a menu-driven interface for performing disk and volume administration functions.

Volume Manager objects created by one interface are fully interoperable and compatible with those created by the other interfaces.

1.2 Disk Array Overview

This section provides an overview of traditional disk arrays.

Performing I/O to disks is a slow process because disks are physical devices that require time to move the heads to the correct position on the disk before reading or writing. If all of the read or write operations are done to individual disks, one at a time, the read-write time can become unmanageable. Performing these operations on multiple disks can help to reduce this problem.

A *disk array* is a collection of disks that appears to the system as one or more virtual disks (also referred to as *volumes*). The virtual disks created by the software controlling the disk array look and act (to the system) like physical disks. Applications that interact with physical disks should work exactly the same with the virtual disks created by the array.

Data is spread across several disks within an array, which allows the disks to share I/O operations. The use of multiple disks for I/O improves I/O performance by increasing the data transfer speed and the overall throughput for the array.

FIGURE 1-23 illustrates a standard disk array.

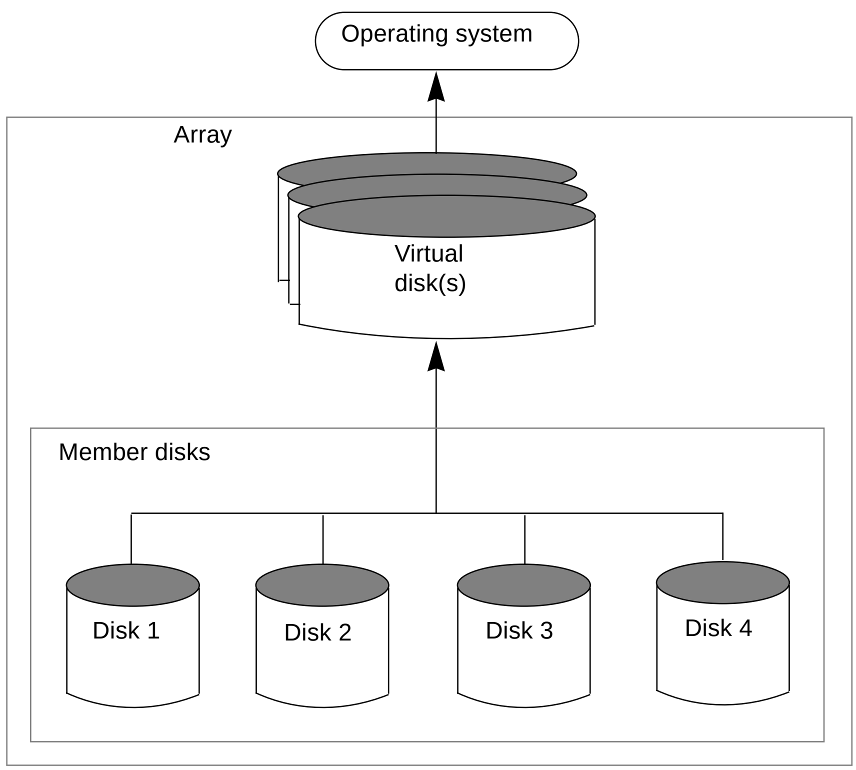

FIGURE 1-23 Standard Disk Array

1.2.1 Redundant Arrays of Inexpensive Disks (RAID)

A *Redundant Array of Inexpensive Disks* (RAID) is a disk array set up so that part of the combined storage capacity is used for storing duplicate information about the data stored in the array. The duplicate information enables you to regenerate the data in case of a disk failure.

Several levels of RAID exist. They are introduced in the following sections.

Note – The Volume Manager supports RAID levels 0, 1, and 5 only.

For information on the Volume Manager's implementations of RAID, refer to Section 1.1.5, "Volume Manager RAID Implementations."

1.2.1.1 RAID-0

Although it does not provide redundancy, striping is often referred to as a form of RAID, known as RAID-0. The Volume Manager's implementation of striping is described in Section 1.1.5.1, "Striping (RAID-0)." RAID-0 offers a high data transfer rate and high I/O throughput, but suffers lower reliability and availability than a single disk.

1.2.1.2 RAID-1

Mirroring is a form of RAID known as RAID-1. The Volume Manager's implementation of mirroring is described in Section 1.1.5.2, "Mirroring (RAID-1)." Mirroring uses equal amounts of disk capacity to store the original plex and its mirror. Everything written to the original plex is also written to any mirrors. RAID-1 provides redundancy of data and offers protection against data loss in the event of physical disk failure.

1.2.1.3 RAID-2

RAID-2 uses bitwise striping across disks and uses additional disks to hold Hamming code check bits. RAID-2 is described in a University of California at Berkeley research paper entitled *A Case for Redundant Arrays of Inexpensive Disks (RAID)*, by David A. Patterson, Garth Gibson, and Randy H. Katz (1987).

RAID-2 deals with error detection, but does not provide error correction. RAID-2 also requires large system block sizes, which limits its use.

1.2.1.4 RAID-3

RAID-3 uses a *parity disk* to provide redundancy. RAID-3 distributes the data in stripes across all but one of the disks in the array. It then writes the parity in the corresponding stripe on the remaining disk. This disk is the parity disk.

FIGURE 1-24 illustrates a RAID-3 disk array.

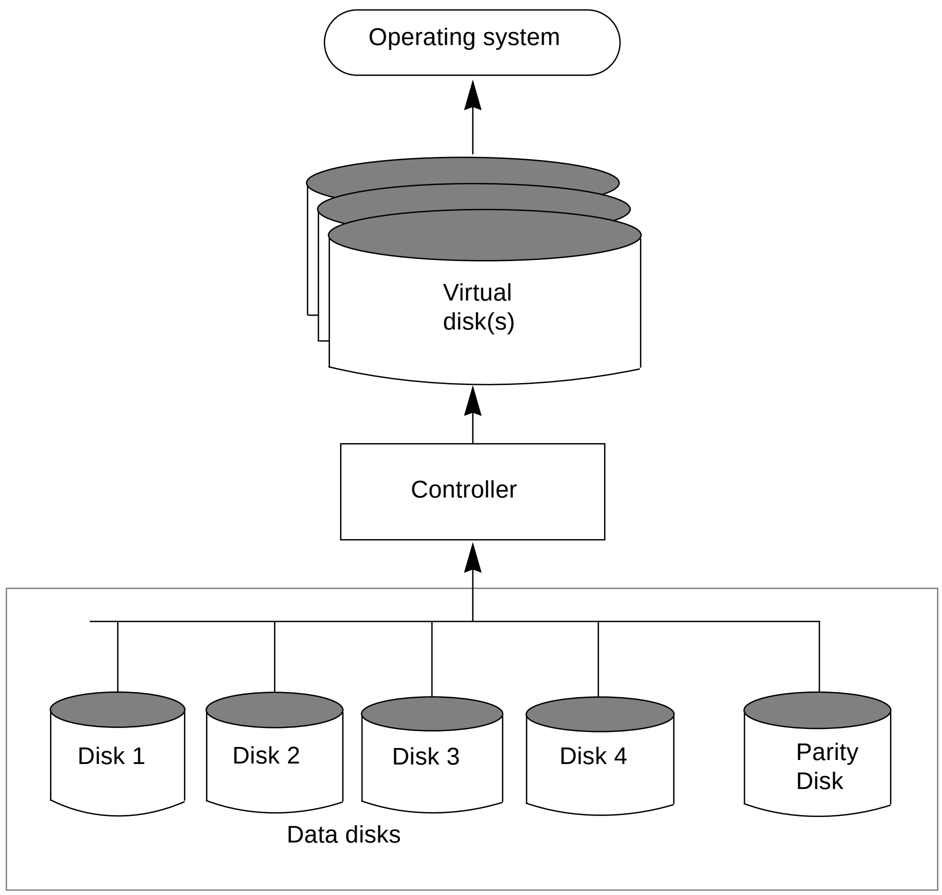

FIGURE 1-24 RAID-3 Disk Array

The user data is striped across the data disks. Each stripe on the parity disk contains the result of an exclusive OR (XOR) procedure done on the data in the data disks. If the data on one of the disks is inaccessible due to hardware or software failure, data can be restored by XORing the contents of the remaining data disks with the parity disk. The data on the failed disk can be rebuilt from the output of the XOR process.

RAID-3 typically uses a very small stripe unit size (also historically known as a *stripe width*), sometimes as small as one byte per disk (which requires special hardware) or one sector (block) per disk.

FIGURE 1-25 illustrates a data write to a RAID-3 array.

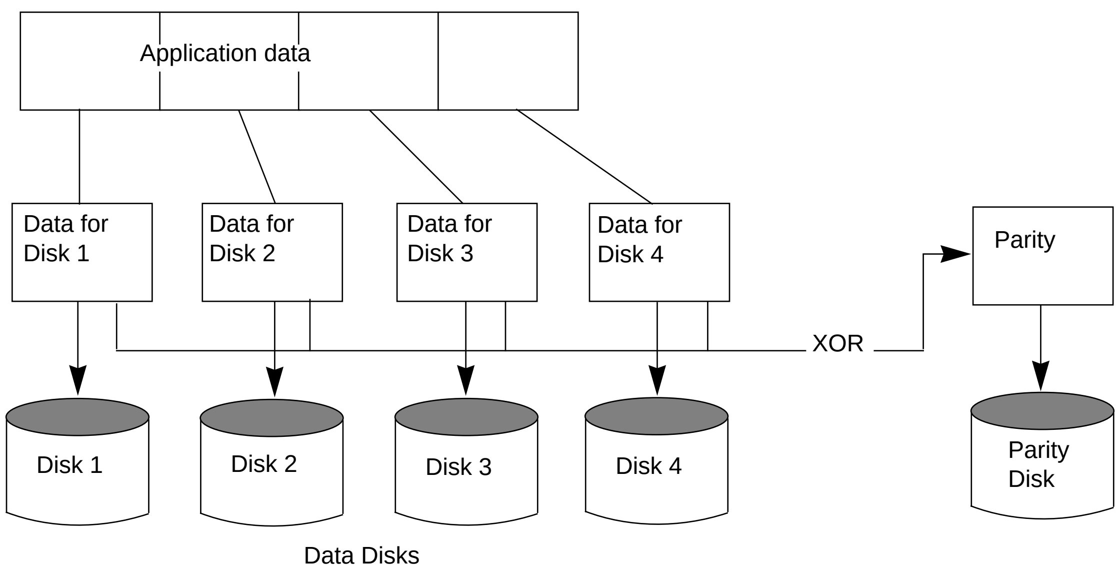

FIGURE 1-25 Data Writes to RAID-3

The parity disk model uses less disk space than mirroring, which uses equal amounts of storage capacity for the original data and the copy.

The RAID-3 model is often used with synchronized spindles in the disk devices. This synchronizes the disk rotation, providing constant rotational delay. This is useful in large parallel writes.

RAID-3 type performance can be emulated by configuring RAID-5 (described later) with very small stripe units.

1.2.1.5 RAID-4

RAID-4 introduces the use of independent-access arrays (also used by RAID-5). With this model, the system does not typically access all disks in the array when executing a single I/O procedure. This is achieved by ensuring that the stripe unit size is sufficiently large that the majority of I/Os to the array will only affect a single disk (for reads).

An array attempts to provide the highest rate of data transfer by spreading the I/O load as evenly as possible across all the disks in the array. In RAID-3, the I/O load is spread across the data disks, as shown in FIGURE 1-25, and each write is executed on all the disks in the array. The data in the data disk is XORed and the parity is written to the parity disk.

RAID-4 maps data and uses parity in the same manner as RAID-3, by striping the data across all the data disks and XORing the data for the information on the parity disk. The difference between RAID-3 and RAID-4 is that RAID-3 accesses all the disks at one time and RAID-4 accesses each disk independently. This enables the RAID-4 array to execute multiple I/O requests simultaneously (provided they map to different member disks), while RAID-3 can only execute one I/O request at a time.

RAID-4's read performance is much higher than its write performance. It performs well with applications requiring high read I/O rates. RAID-4 performance is not as high in small, write-intensive applications.

The parity disk can cause a bottleneck in the performance of RAID-4. This is because all the writes that are taking place simultaneously on the data disks must each wait its turn to write to the parity disk. The transfer rate of the entire RAID-4 array in a write-intensive application is limited to the transfer rate of the parity disk.

Since RAID-4 is limited to parity on one disk only, it is less useful than RAID-5.

1.2.1.6 RAID-5

RAID-5 is similar to RAID-4, using striping to spread the data over all the disks in the array and using independent access. However, RAID-5 differs from RAID-4 in that the parity is striped across all the disks in the array, rather than being concentrated on a single parity disk. This breaks the write bottleneck caused by the single parity disk write in the RAID-4 model.

FIGURE 1-26 illustrates parity locations in a RAID-5 array configuration. Every stripe has a column containing a parity stripe unit and columns containing data. The parity is spread over all of the disks in the array, reducing the write time for large independent writes because the writes do not have to wait until a single parity disk can accept the data.

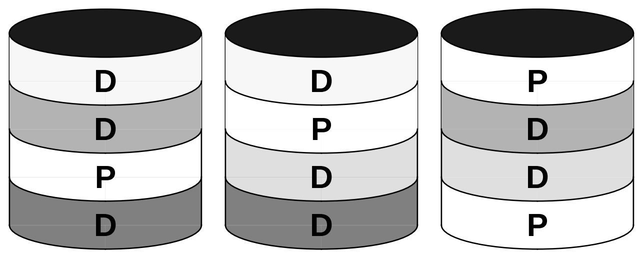

D = Data stripe unit
P = Parity stripe unit

FIGURE 1-26 Parity Locations in a RAID-5 Model

For additional information on RAID-5 and how it is implemented by the Volume Manager, refer to Section 1.1.5.4, “Volume Manager and RAID-5.”

1.3 Dynamic Multipathing

The Volume Manager actively supports multiported disk arrays. It automatically recognizes multiple I/O paths to a particular disk device within the disk array. The Dynamic Multipathing feature of the Volume Manager provides greater reliability by providing a path failover mechanism. In the event of a loss of one connection to a disk, the system continues to access the critical data over the other healthy connections to the disk. The multipathing functionality also provides greater I/O throughput by balancing the I/O load uniformly across multiple I/O paths to the disk device.

In the Volume Manager, all the physical disks connected to the system are represented as metadevices with one or more physical access paths. A single physical disk connected to the system is represented by a metadevice with one path, whereas a disk that is part of a disk array such as Sun Enterprise Network Array, is represented by a metadevice that has two physical access paths. The user can use the Volume Manager administrative utilities such as `vxdisk` to display all the paths of a metadevice, and the status information of the various paths.

CHAPTER 2

Introduction to the Visual Administrator

This chapter describes the graphical user interface for the Volume Manager, called the *Visual Administrator* (vxva).

The quick reference sections at the end of this chapter provide information on menus and forms associated with Visual Administrator windows.

Note – You should understand Volume Manager objects before you use the Visual Administrator. Refer to Chapter 1 for more information.

2.1 Purpose of the Visual Administrator

The Visual Administrator's primary function is to provide a graphical user interface to the Volume Manager. In addition, the Visual Administrator acts as an interface to several common file system operations; some of these file system operations are supported for the VERITAS File System (VxFS) only.

The Visual Administrator interface is relatively intuitive and easy to use. Windows, menus, forms, and icons are provided to help you. The Visual Administrator is one of several interfaces to the Volume Manager. Volume Manager objects created by one interface are fully interoperable and compatible with those created by the other interfaces.

2.2 Visual Administrator Icons

The Visual Administrator represents the various Volume Manager objects through icons. When a change is made to existing Volume Manager objects using the Visual Administrator (or another Volume Manger interface), every open Visual Administrator session on the same system automatically adjusts its icons to reflect the change.

Volumes are usually composed of plexes (mirrors), which are composed of subdisks. Volume icons therefore typically contain associated plex icons which, in turn, contain associated subdisk icons. While both volume and plex icons can appear alone and with no associated components, they are not useful in this form.

The Visual Administrator uses the following icons to represent Volume Manager objects.

Physical Disk Icon

Physical disks appear as cylindrical icons labeled PD. These icons contain rectangular partition icons. Physical disk icons represent physical disks known to the Visual Administrator. Physical disk icons appear in the View of Disks.

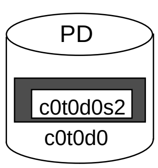

FIGURE 2-1 Physical Disk Icon

Partition Icon

Partitions appear as rectangular icons within physical disk icons. A partition icon is labeled with the name of the physical disk with which it is associated (disk access name). If a VM disk exists for a partition (and its corresponding VM disk icon appears in a disk group view), the partition icon is shaded. A partition usually represents slice 2 of the disk (all application-accessible space on the disk). Partition icons appear in the View of Disks.

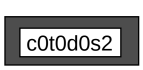

FIGURE 2-2 Partition Icon

VM Disk Icon

VM disks appear as cylindrical icons labeled `D`. They usually contain subdisks, which are represented as rectangles. A VM disk is one that is both under Volume Manager control and assigned to a disk group. VM disk icons represent the disk media records used to map the logical aspect of VM disks to physical disks. A VM disk icon is labeled with its disk media name. The name of the partition with which the VM disk is associated appears below the VM disk icon. VM disk icons typically appear in a disk group view. If a VM disk is designated as a hot-relocation spare, its icon label changes from `D` to `S`.

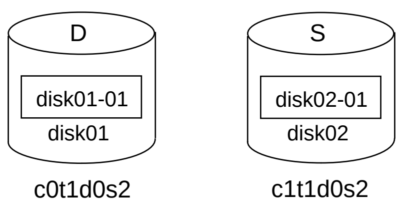

FIGURE 2-3 VM Disk Icons

Subdisk Icon

Subdisks appear within VM disks as rectangular shaped icons. Subdisk icons are also duplicated within any plex icons with which they are associated. In plex and VM disk icons, subdisk icons that represent non-contiguous storage are visually separated by small gaps. Subdisk icons typically appear in disk group views and in the View of Volumes.

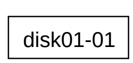

FIGURE 2-4 Subdisk Icon

Log subdisks (used to log disk activity for Dirty Region Logging or RAID-5 logs) have icons with double borders to distinguish them from regular subdisk icons.

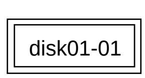

FIGURE 2-5 Log Subdisk Icon

Plex Icon

Plexes appear either alone or within volumes as relatively large rectangles containing subdisks. Plex icons have a heavy border to distinguish them from partition or subdisk icons. Plex icons typically appear in disk group views and in the View of Volumes.

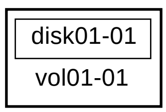

FIGURE 2-6 Plex Icon

Striped and concatenated plexes differ visually in that a striped plex has gaps between its subdisks, while the subdisks of a concatenated plex are joined. FIGURE 2-7 illustrates the difference between concatenated (`vol01-01`) and striped (`vol02-01`) plexes.

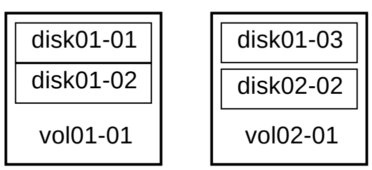

FIGURE 2-7 Concatenated and Striped Plex Icons

In its simplest form, a striped plex has one subdisk per stripe column. However, multiple subdisks can exist per stripe column.

FIGURE 2-8 illustrates a striped plex with one subdisk in the top stripe column and multiple subdisks in the middle and bottom columns. The bottom column contains a gap where a subdisk has been removed; this gap can be filled with one or more subdisks.

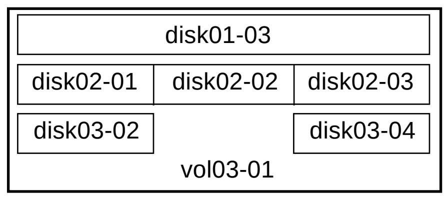

FIGURE 2-8 Striped Plex With Multiple Subdisks per Column

Volume Icon

Volumes appear as cylindrical icons labeled V. These icons often contain plex and subdisk icons. Volume icons are distinguished from disk icons by a heavy border. Volume icons typically appear in disk group views and in the View of Volumes.

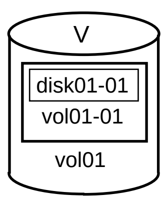

FIGURE 2-9 Volume Icon

Mirrored volumes contain two or more plex icons. The following is a mirrored volume with one striped plex (`vol01-01`) and one concatenated plex (`vol01-02`).

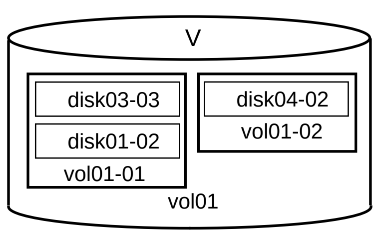

FIGURE 2-10 Mirrored Volume Icon

A RAID-5 volume contains a single RAID-5 plex icon and may contain one or more log plex icons. The following is a RAID-5 volume with one RAID-5 plex (`vol01-01`) and one log plex (`vol01-02`). RAID-5 volumes are labeled `RAID-5` instead of V.

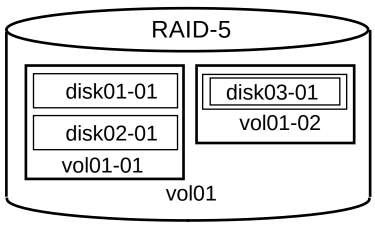

FIGURE 2-11 RAID-5 Volume Icon

2.2.1 Busy Icons

With some operations, icons are updated almost instantly to reflect the results of the operation just performed. During other operations, it may take awhile for a particular icon to update itself. While being updated, icons are prevented from accepting input or undergoing configuration changes. The Visual Administrator dims the text in that icon so that you are aware that it is temporarily unavailable for further configuration operations. No input is accepted by an icon while it is dimmed. As soon as the icon is fully updated, it returns to its normal visual state and accepts input again. Icons that are temporarily dimmed in this manner are also referred to as *blocked icons*.

FIGURE 2-12 illustrates a plex icon that is temporarily dimmed.

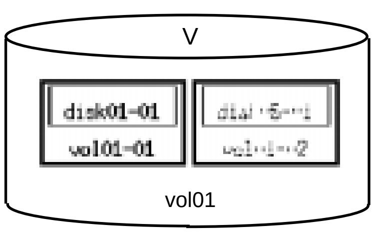

FIGURE 2-12 Busy Icon

2.3 Using the Mouse With the Visual Administrator

You must have a two-button or three-button mouse to use the Visual Administrator. The Visual Administrator enables you to use the mouse (either alone or together with certain keyboard keys) to select, manipulate, or display the properties of icons and their underlying objects.

TABLE 2-1 describes terms associated with the use of the mouse.

TABLE 2-1

Term	Definition
Click	Press and release the mouse button in a single, smooth motion.
Double-click	Click the mouse button twice in rapid succession.
Press	Press and continue to hold down the mouse button.
Point	Move the tip of the pointer onto an item on the screen.
Drag and drop	Press and continue to hold down the mouse button while moving the mouse and pointer (to drag an object). Release the button when the pointer reaches its destination (to drop the object).
Select	Click the mouse button while the pointer is directly over the item to be selected (such as an icon).

Throughout this guide, the mouse buttons are referred to as the LEFT, MIDDLE, and RIGHT buttons. TABLE 2-2 describes the default mouse buttons and their locations, which vary according to the number of buttons on the mouse.

TABLE 2-2

Virtual Mouse Button	3-Button Access	2-Button Access
LEFT	LEFT	LEFT
MIDDLE	MIDDLE	Ctrl-LEFT
RIGHT	RIGHT	RIGHT
Shift-LEFT	Shift-LEFT	Shift-LEFT
Shift-MIDDLE	Shift-MIDDLE	Ctrl-RIGHT
Shift-RIGHT	Shift-RIGHT	Shift-RIGHT

Note – This guide assumes you are using a 3-button, right-handed mouse. Button 1 is the LEFT button and Button 3 is the RIGHT button. It is possible to remap mouse buttons (using the `xmodmap` command, for instance). Refer to your X Window System documentation for details.

TABLE 2-3 lists the mouse and keyboard combinations recognized by the Visual Administrator.

TABLE 2-3

Action	Result
LEFT	Selects a single icon. Clicking LEFT on an icon causes that icon to be the currently-selected one; any previously-selected icons are de-selected.
Shift-LEFT	Toggles between minimizing or maximizing an icon. A maximized icon displays all of its components in detail. A minimized icon is compressed and its components are concealed.
MIDDLE	Selects either one icon or multiple icons simultaneously. When selecting with the MIDDLE button, any previously-selected icons remain selected along with the newly-selected one. Clicking MIDDLE on a selected icon de-selects that icon.
Shift-MIDDLE	Toggles between starting or stopping projection on the selected icon.
RIGHT	On an icon that is *not* undergoing analysis, displays the properties form for that object. On an icon that is undergoing analysis, displays the analysis statistics form for that object.
Shift-RIGHT	Always displays the properties form for the object, regardless of whether analysis is in effect.

2.4 Views

Views are special Visual Administrator windows that display icons representing all or a subset of the objects currently known to the Visual Administrator. With views, you can examine and manipulate different parts of the physical and logical storage systems. Icons can be added to or removed from views only with the Visual Administrator.

In a view window, you can use the following menu features:

TABLE 2-4

Menu	Purpose
File	Close the window or exit the Visual Administrator completely
Basic-Ops	Perform file system, volume, and disk operations
Advanced-Ops	Perform volume, plex, subdisk, disk group, and disk operations
Analyze	Display visual or numerical statistics for a Volume Manager object's performance
Projection	Display visual information about the properties and associations between icons
Options	Set user preferences
Icon	Perform assorted operations on icons
Help	Open a help facility window

Each view window title includes the name of the machine on which the Visual Administrator session is running. If the Visual Administrator is running in demo mode, the string `(DEMO)` appears in the view window title instead.

Once the Visual Administrator is started, the *Visual Administrator main window* appears. This window provides a view button area (the *views subwindow*) containing a button for every view on the system. Views are accessed by clicking the LEFT mouse button on the appropriate view button in the views subwindow.

Each *disk group* (a collection of disks that share a common configuration) is represented by the Visual Administrator as a view window instead of an icon. The VM disks, volumes, and other objects belonging to a particular disk group are all displayed within that disk group's view.

FIGURE 2-13 illustrates the Visual Administrator main window.

FIGURE 2-13 Visual Administrator Main Window

2.4.1 Types of Views

There are two types of views in the Visual Administrator: default and user-created views. Both view types function identically, but certain restrictions are placed on default views. Default views cannot be removed or renamed by users, unlike user-created views.

2.4.1.1 Default Views

The default views provided by the Visual Administrator include:

- View of `rootdg` — Displays all objects that have been assigned to the default disk group, `rootdg` (including Volume Manager disks, volumes, and associated objects). If additional disk groups are created, the Visual Administrator creates views for them. Disk group views can be used for most operations.
- View of Disks — Displays all physical disks on the system showing an icon for each physical disk. Each disk that has been initialized for Volume Manager use has a device entry on the system. This view is an iconic representation of the physical storage subsystem. This view is used for disk and disk group operations only.
- View of Volumes — Displays all volumes (as well as plexes and associated subdisks) on the system. It displays all Volume Manager volumes, regardless of disk group. Plexes and any associated subdisks are also displayed here. Although this view can be used for various operations, it is preferable to use a disk group view for volume operations.
- View of World — Displays all objects on the system (including physical disks, volumes, and associated objects). It displays all Volume Manager disks, volumes, and plexes, regardless of disk group. This view essentially combines both the Disks and Volumes views. It also contains any unassociated plexes, which do not

appear in most other views. Although this view can be used for various operations, it is generally preferable to use a disk group view or the Disks view for disk operations.

FIGURE 2-14 illustrates a View of `rootdg` window that does not yet contain any volumes.

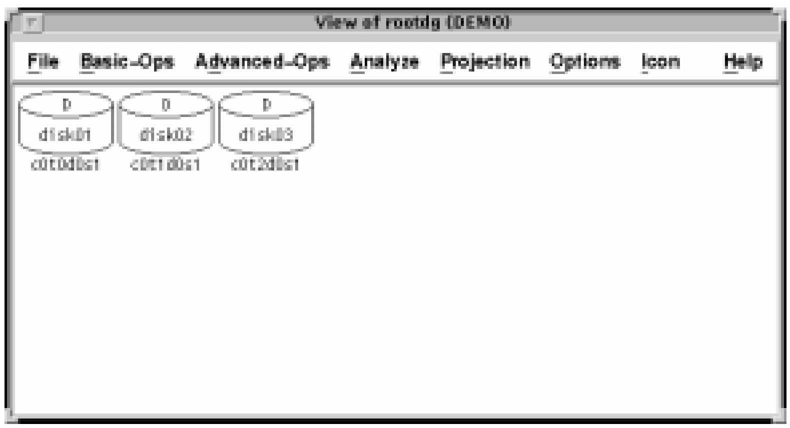

FIGURE 2-14 View of `rootdg`

2.4.1.2 User-Created Views

A user-created view is a view window that focuses on a particular part of a physical and/or logical storage system, as defined by the system administrator. The system administrator can create views consisting of a selected collection of icons. For example, you can create a special view to correspond to a physical or logical grouping (such as a view for the accounting department). User-created views enable you to isolate part of the storage subsystem in order to observe or monitor that part of the configuration.

User-created views differ from default views in that they contain *copies* of icons from default views. Operations performed on these icon copies are reflected in the default views that display the affected icon(s).

User-created views can be created using the Views pull-down menu from the Visual Administrator main window. Once created, icons can be added to a new view window by copying them over from existing views via the Icon menu, or by dragging them into the desired view.

For detailed information on creating, renaming, or removing views, refer to TABLE 2-6, TABLE 2-7, and Section 2.7.2.3, "Remove View Form."

2.5 Color and Bitmap Patterns

Depending on the type of monitor being used, Visual Administrator uses color or bitmap patterns to indicate:

- The state of an icon
- The activity level of an icon
- Relationships between icons
- Failure of an operation

If a color monitor is used, the default colors are red, yellow, grey, and green. If a monochrome monitor is used, bitmap patterns of varying textures and shades are used instead of colors. By default, standard X Window System bitmaps (typically located in `/usr/openwin/include/X11/bitmaps`) are used to create these patterns.

TABLE 2-5 gives the values for the default colors and bitmap patterns associated with Visual Administrator icons under different conditions.

TABLE 2-5

Situation	Color	Bitmap Pattern
Selected icon	Green	grey3
Disabled icon	Grey	stripe4
Alarmed icon	Red	grey1
Free subdisk icon	Yellow	root_weave
Projection	Yellow	root_weave
Analysis: low	Green	cross_weave
Analysis: medium	Yellow	root_weave
Analysis: high	Red	wide_weave

Note – This guide assumes that the default colors or bitmaps are being used. However, colors or bitmaps (as well as other interface preferences) can be redefined in your `.Xdefaults` file. Refer to Appendix A for information on resetting these defaults.

It is possible for a single icon to be in multiple states represented by different colors or patterns at once. For example, a given icon may be both selected and under projection at the same time. In such cases, the Visual Administrator reflects the color or pattern that represents the highest priority. The following is the priority list for possible icon states, starting with the highest priority:

1. Blocked
2. Error
3. Selected
4. Projected
5. Analyzed
6. Enabled

An icon that is in the blocked state (highest priority) is one that is currently busy and cannot allow any mouse or keyboard input. The text within a blocked icon is dimmed to indicate that it is inaccessible.

2.6 Basic Versus Advanced Operations

The Visual Administrator can be used to perform several types of volume management and other administrative operations. Most of these operations can be invoked from the Basic-Ops and Advanced-Ops menus in the menu bars of the views windows.

The basic and advanced operations menus provide two distinct approaches to volume management:

- Basic — The basic approach typically uses the Volume Manager’s `vxassist` utility, which takes information about what you want to accomplish and automatically performs the necessary underlying tasks. This approach requires only minimal input, but does permit more detailed specifications. Basic operations are accessed via the Basic-Ops menu.
- Advanced — The advanced approach typically uses the Volume Manger command set, which consists of a number of fairly complicated commands that often require you to specify detailed input. The Volume Manger commands use a “building block” approach that requires you to have a detailed knowledge of the underlying Volume Manger structure and components to manually perform the sequences of commands necessary to accomplish a certain task. Advanced operations are accessed via the Advanced-Ops menu.

Some operations are available through one approach only, while others can be handled through either the Basic-Ops or Advanced-Ops menu.

Note – When creating and manipulating volumes, first select which approach to use and then proceed with that approach only. The basic and advanced approaches cannot be combined to perform a single volume operation. Use the basic approach whenever possible.

2.7 Quick Reference to Views Menus and Forms

The Visual Administrator main window contains a Views menu, which you use to create and modify view windows.

2.7.1 Views Menus

You can access the following options via the Views menu (located in the Visual Administrator main window only).

TABLE 2-6

Menu /Submenu Access	Description
Views ➤ Create a View	Creates a new user-specified view (also known as a *user-created view* or *user view*). Specify a unique name for this new view. Once created, icons can be added to this new view window by copying them over from existing views. User-created views differ from default views (disk, volume, disk group, and world views) in that they contain copies of icons from the default views. **See Also**: * View Create Form (TABLE 2-7)
Views ➤ Rename a View	Renames an existing user-created view. Specify a new, unique name for the view. **Requirements**: * Only a user-created view can be renamed. * The new view name must be unique. **See Also**: * Rename View Form (TABLE 2-8)
Views ➤ Remove Views	Permanently removes a user-created view. The removal of a user-created view results in the removal of its icons, but the objects represented by those icons are unaffected. Requirements: * Only a user-created view can be removed. **See Also**: * Remove View Form (Section 2.7.2.3, “Remove View Form”)

2.7.2 Views Forms

You can access the following forms via the Views menu (located in the Visual Administrator main window only).

2.7.2.1 View Create Form

This form is used to create a new view window. Once created, this view is represented by a new view button in the Visual Administrator main window. This user-created view is also retained across future Visual Administrator sessions.

TABLE 2-7

Field	Description
View Name:	Determines the name of the view to be created. This name must be unique. The maximum length of the name is 31 characters. This field is required. You can delete or rename this view at a later time.

2.7.2.2 Rename View Form

This form is used to change the name of an existing user-created view.

TABLE 2-8

Field	Description
Old View Name:	Lists the name of the view to be renamed. This name must belong to a user-created view. This field is required.
New View Name:	Lists the name to which the existing name is to be changed. This name must be unique. The maximum length of the name is 31 characters. This field is required.

2.7.2.3 Remove View Form

Use this form to remove a user-created view. *This form contains no fields*. Instead, it lists all user-created views that currently exist.

1. **Select one or more views for removal by clicking on the desired view name(s) to highlight it.**

   If you mistakenly select a wrong view, click on it again to clear that view name.

2. **Click Apply, then click Apply again on the confirmation window.**

CHAPTER **3**

Getting Started

This chapter describes how to start the Visual Administrator and perform some basic operations using this graphical user interface.

The quick reference sections at the end of this chapter describe Visual Administrator menus and forms related to the topics discussed in this chapter.

3.1 Initializing the Visual Administrator

Before the Visual Administrator can be started, both the Volume Manager and the Visual Administrator must be properly installed. In addition, the Volume Manager configuration daemon (`vxconfigd`) must be running and enabled. If you have followed the procedures outlined in the *Sun StorEdge Volume Manager 2.6 Installation Guide*, these conditions should already be met.

In order to run the Visual Administrator from the command line, your `PATH` environment variable must contain `/opt/SUNWvxva/bin`. If necessary, update `PATH` to include `/opt/SUNWvxva/bin` before attempting to run the Visual Administrator.

- If using a Bourne Shell (`sh` or `ksh`), update `PATH` as follows:

```
PATH=${PATH}:/opt/SUNWvxva/bin
export PATH
```

- If using a C Shell (`csh` or `tcsh`), update `PATH` as follows:

```
setenv PATH ${PATH}:/opt/SUNWvxva/bin
```

3.1.1 Demo Mode

A demo mode is provided so that you can experiment with the Visual Administrator without using a real system. Demo mode sets up a series of dummy disks on which most Visual Administrator operations can be performed. However, file system operations and basic disk operations are not accessible in demo mode and are therefore dimmed in the Basic-Ops menu.

In demo mode, the view window titles include the string `(DEMO)`. In real mode, the view window titles include the string (*hostname*).

User-created configurations are maintained between demo-mode sessions. However, you can run the `/opt/SUNWvxva/bin/vxva_setup` script at any time to clean up demo mode configurations and return to the default demo configuration.

3.1.2 Command Line Startup

`vxva` can be run in real mode or in demo mode.

Note – Only users with superuser privileges can run the Visual Administrator in real mode. Demo mode requires no special privileges.

Start the Visual Administrator from the command line in either of the following ways:

▼ To Run the Visual Administrator in Real Mode

- **Type the following line at the prompt:**

```
# vxva &
```

▼ To Run the Visual Administrator in Demo Mode

- **Type the following line at the prompt:**

```
vxva -t &
```

A temporary window is displayed stating that the Visual Administrator is starting up.

3.1.3 The Visual Administrator Main Window

The Visual Administrator main window is displayed immediately after the Visual Administrator is properly invoked (see FIGURE 3-1).

FIGURE 3-1 Visual Administrator Main Window

The Visual Administrator main window provides access to view windows and allows you to do any of the following:

- Exit the Visual Administrator through the File menu
- Add, rename, or remove user-created views through the Views menu
- Set user preferences through the Options menu
- Open a help facility window through the Help menu
- Open the View of Disks window with the Disks view button
- Open the View of Volumes window with the Volumes view button
- Open the View of World window with the World view button
- Open the View of `rootdg` window with the `rootdg` view button

From any view window, the Basic-Ops or Advanced-Ops menu can be used to create new Volume Manager objects (represented by icons) or to manipulate existing icons.

3.2 Performing Basic Operations

The Visual Administrator basic operations enable you to perform administrative operations relatively quickly and easily. The basic operations use free space management techniques to automatically locate available space for the desired operation, unless you specify some suitable disk space.

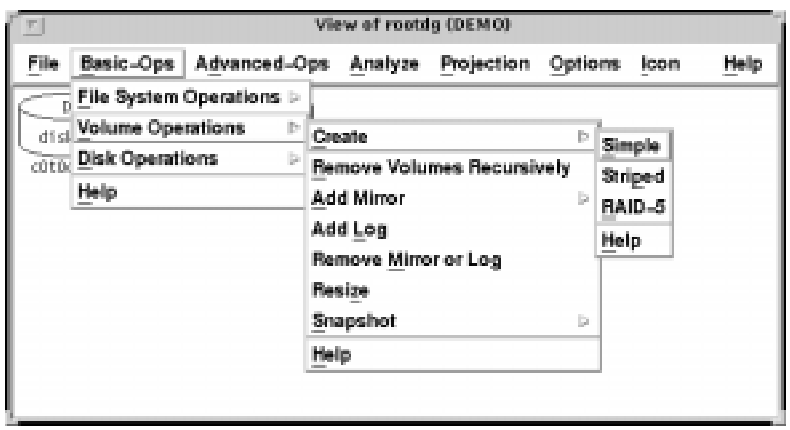

FIGURE 3-2 Basic-Ops Menu

The Basic-Ops menu provides submenus (FIGURE 3-2) for the following operations:

- File System Operations — For general file system maintenance. File systems can be created, mounted and unmounted, resized, and backed up.
- Volume Operations — For general volume maintenance. Volumes can be created from Volume Manager disks, removed, mirrored, resized, and backed up.
- Disk Operations — For general disk maintenance. Disks can be initialized, removed, or replaced.

Note – File system operations and basic disk operations are not simulated by the Visual Administrator in demo mode; they are only accessible in real mode. The advanced disk operations are, however, accessible in demo mode.

3.3 Performing Advanced Operations

The Visual Administrator advanced operations enable you to manipulate Volume Manager entities manually. You must use the Advanced-Ops menu for those operations that are not available via the Basic-Ops menu.

Be sure you have a solid understanding of Volume Manager concepts before you try to create and manipulate volumes using the Advanced-Ops menu.

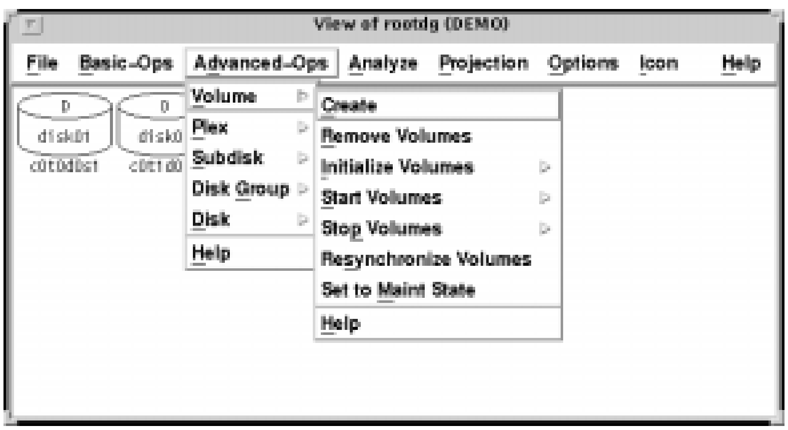

FIGURE 3-3 Advanced-Ops Menu

The Advanced-Ops menu provides submenus (FIGURE 3-3) for the following operations:

- Volume — Perform operations on one or more volumes. Possible operations include creating a new volume, removing a volume, and changing the state of an existing volume.
- Plex — Perform operations on one or more plexes (mirrors). Possible operations include creating a new plex, removing a plex, associating a plex with a volume, disassociating a plex from a volume, and detaching a plex.
- Subdisk — Perform operations on one or more subdisks. Possible operations include creating a new subdisk, removing a subdisk, associating a subdisk with a plex, disassociating a subdisk from a plex, and joining or splitting subdisks.
- Disk Group — Perform operations on disk groups. Possible operations include creating a new disk group, adding disks to a disk group, and removing disks from a disk group.

- Disk — Perform operations on disks. Possible operations include initializing a disk, removing a disk from VxVM control, and onlining or offlining a disk.

3.4 Checking Free Space

For operations that require free disk space, the Volume Manager and Visual Administrator usually allow you to designate disks with sufficient free space. If no disks are specified, disks with available space are automatically used when an operation is performed via the Basic-Ops menu. With Advanced-Ops, however, you are expected to designate disks with sufficient space.

Before designating a disk to be used, verify the amount of free space on that disk.

▼ To Check the Available Free Space on a Disk

- **Click the RIGHT button on a portion of the VM disk icon that is not covered by another icon.**

The VM disk properties form is displayed. A field at the bottom of the form shows the maximum free space available on this disk. For detailed information on the VM disk properties form, see TABLE 4-11.

3.5 Displaying Properties of an Object

The properties of the object corresponding to a given icon can be displayed (and potentially adjusted) in a properties form. Properties forms provide detailed information about the characteristics of a particular object.

Properties forms exist for:

- Physical disks
- Partitions
- VM disks
- Volumes
- Plexes
- Subdisks
- File systems

Since file systems are not represented by icons, the file system properties form is accessed through the File System Operations menu. If the file system is mounted, the file system properties form can be accessed by clicking RIGHT on the mount point name that appears below its volume icon.

▼ To Access a Properties Form for an Icon

- **Move the pointer over the icon and click the RIGHT button (if the icon is under analysis, Shift-RIGHT must be used instead).**

  When the properties form is displayed, details of the icon's configuration are displayed. You can alter some of the displayed properties directly through this form by altering the appropriate fields and then clicking the form's Apply button. Clicking the Reset button converts any altered form fields back to their values at the time when the form was popped up.

  FIGURE 3-4 shows a typical properties form.

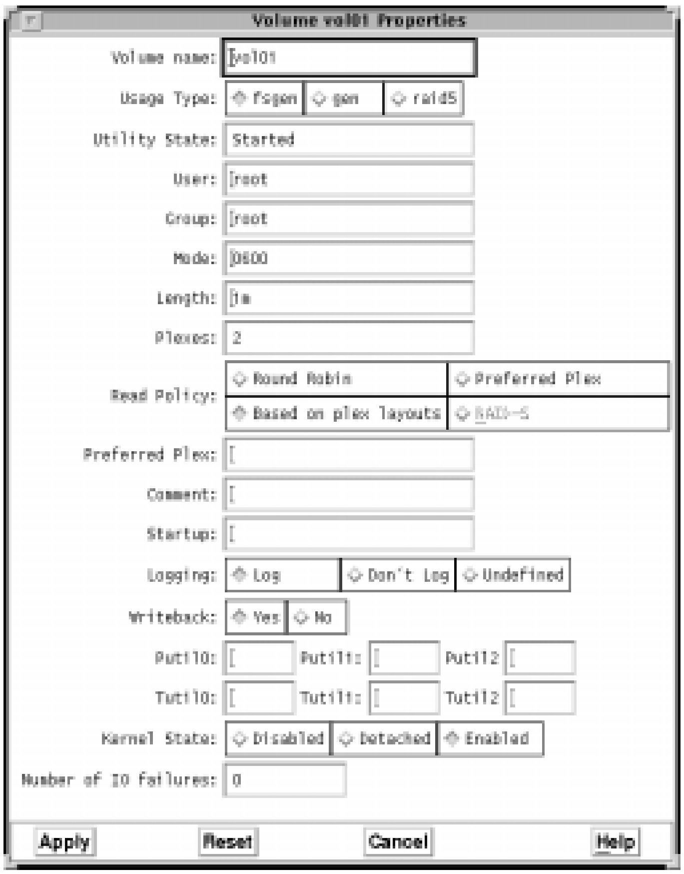

FIGURE 3-4 Properties Form

3.5.1 To Rename an Object

1. **Click RIGHT on the icon corresponding to the object to be renamed (if the icon is under analysis, use Shift-RIGHT instead.**

   The object's properties form is displayed.

2. **Alter the name field for that object.**

3. **Click Apply.**

   The icon corresponding to the renamed object will adjust its name accordingly.

Note – Be sure that no application is accessing the object by its old name when you are renaming it.

3.6 Invoking Projection on Objects

Projection is the technique the Visual Administrator uses to show relationships between icons that represent Volume Manager objects. Projection is indicated using color (yellow is the default) or bitmap patterns. Projection highlights those objects that the selected object is composed of and illustrates the relationship between the objects. For example, if a volume is selected for projection, the corresponding subdisks are highlighted within the volume icon and also on the appropriate disk icons. If the selected icon has no associated objects, the Visual Administrator issues a warning to this effect.

▼ To Show the Projection for a Particular Icon

- **Click MIDDLE on the icon while holding down the Shift key (Shift-MIDDLE).**

▼ To Stop Projection of a Particular Icon

- **Click Shift-MIDDLE on the icon again.**

You can also start or stop projection by selecting an icon and then using the Icon Projection submenu of the Projection menu. For detailed information on projection menus, see TABLE 3-4.

Projection can be requested in any view. When an icon is highlighted by projection, all icons representing that object in all view windows where it appears are highlighted.

TABLE 3-1 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 3-1

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.
VM Disk	All plexes associated with the subdisks that reside on the disk.
Physical Disk	All associated partition and VM disk icons.
Partition	The VM disk associated with the partition.

When projection is turned on and left on from two different objects, any icon that happens to be related to both of these objects receives two “layers” of projection highlighting. Projection must then be turned off from both objects that started it (or all projection in the current session must be stopped) before the double-highlighted icon returns to its normal state. For example, if projection is turned on from both a volume and a plex related to the same subdisk, then that subdisk is highlighted twice even though it appears to have only one layer of highlighting.

FIGURE 3-5 illustrates highlighting that results from the selection of a volume icon (`vol01`) for projection.

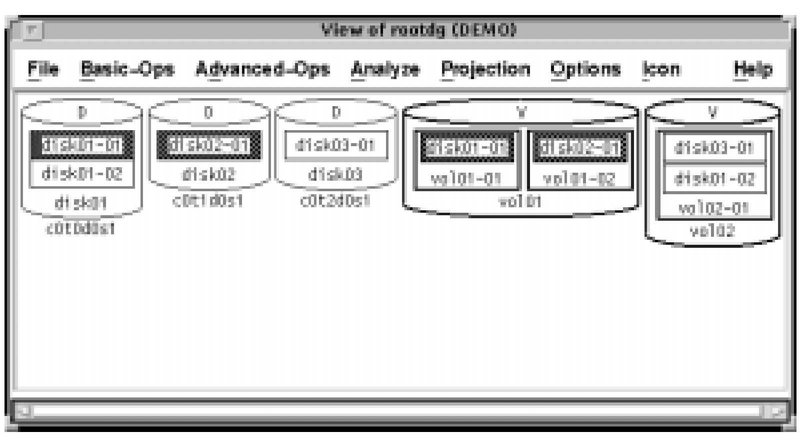

FIGURE 3-5 Projection

3.7 Displaying Activity Levels on Objects

Through analysis, the Visual Administrator displays statistics about the performance of Volume Manager objects. These statistics are displayed both visually (with different colors or patterns) and numerically (with pop-up statistics forms). Analysis displays activity levels on selected icons representing Volume Manager objects. Analysis statistics can be written to a log file using the Analysis Parameters form (for detailed information on this form see TABLE 3-6). When a log file is specified, time-stamped statistics are written to the file at each interval in binary format. Since statistics can only be collected on volumes, only subdisks that are associated with volumes, and disks associated with subdisks associated with volumes, can be analyzed. For detailed information on the Analysis Statistics form, see TABLE 3-7.

Note – Only volume and VM disk icons can be selected for analysis.

▼ To Start Analysis

1. **Select one or more VM disk and/or volume icons.**
2. **Choose Start from the Analyze menu.**

   Volumes, VM disks, and subdisks that are associated with the selected icon(s) change their colors or patterns to reflect their relative activity levels (high, medium, or low). FIGURE 3-6 shows typical analysis for a selected volume (`vol01`).

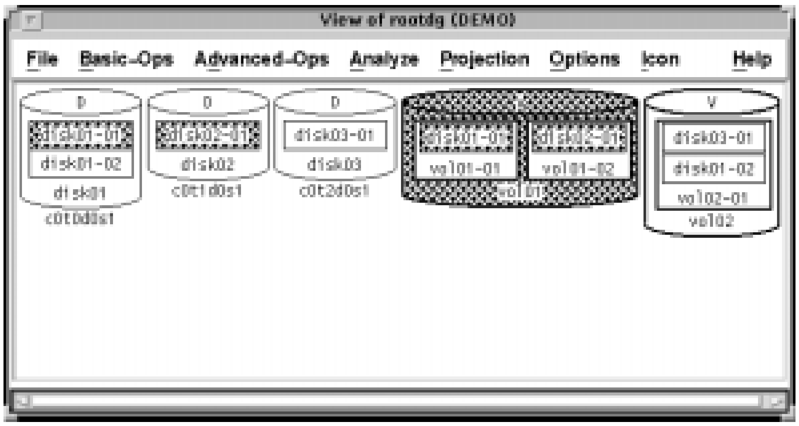

FIGURE 3-6 Analysis

▼ To Stop Analysis

1. **Select the icons on which analysis is to stop.**
2. **Choose Stop from the Analyze menu.**

3.7.1 Analysis Table

On a color display, the icons being analyzed are colored to indicate their statistics. On a monochrome display, analysis is depicted with bitmap patterns instead of colors.

Using the Analysis Parameters form (TABLE 3-6), you can specify how you want the statistics to correspond to Low, Medium, and High ranges. The default colors and patterns can also be changed using the X resources for the Visual Administrator. See Appendix A, “X Resources for the Visual Administrator” for information on these resources.

TABLE 3-2 summarizes the default colors and patterns associated with the various levels of analysis.

TABLE 3-2

Analysis Level	Color	Bitmap Pattern
Low	Green	cross_weave
Medium	Yellow	root_weave
High	Red	wide_weave

3.8 Changing Options

Most Visual Administrator windows contain an Options menu, which enables you to specify preferences for Visual Administrator operations.

Some of the preferences have associated submenus that require further selection.

User preferences are saved in the `$HOME/.vxva_pref` file so that they are retained across Visual Administrator sessions.

The preferences that can be set through the Options menu are:

- Show Command
- When Commands are Ready
- Logging
- Popup the Command Window
- Format of Size
- Help

Some of the available preference settings relate to the Command Info Window. This is a special window that displays the command history, along with the status and results of those commands.

3.8.1 Show Command

This specifies whether the Command Info Window is to be shown before every command is executed (Show at Start), or only when a command fails (Show on Error). If no preference is indicated, the default is Show on Error.

3.8.2 When Commands Are Ready

This specifies what the Visual Administrator should do when it is ready to run a Volume Manager or system command. The command can either be run immediately (Execute Commands), or brought up in the Command Info Window for inspection (Show Commands Only). If no preference is indicated, the default is Execute Commands.

If Show Commands Only is selected, the command corresponding to a Visual Administrator operation is displayed but it is not executed. No changes are made to the existing configuration and analysis is unaffected. This option makes it possible to preview the VxVM or system command underlying a specific Visual Administrator operation. The command can then be executed directly from the Command Info Window once it is approved.

3.8.3 Logging

This indicates whether command logging should be used. Logging keeps a record on disk of all commands sent to the Volume Manager or the system by the Visual Administrator. Logging can be started (Start) or stopped (Stop) at any time. When started, a window requesting a log file name appears and a file name must be entered; the Apply button must then be selected to start logging with the designated file. If no preference is indicated, the default is Stop and no logging is in effect. For detailed information on the Log File form, refer to TABLE 3-9.

Note – Unlike other user preferences, the logging setting is not saved across Visual Administrator sessions.

3.8.4 Popup the Command Window

This option opens the Command Info Window. This window displays the command history, along with the status and results of those commands. Commands can also be executed or repeated from this window. The Command Info Window is described in more detail in Section 3.9.2, “Using the Command Info Window”.

3.8.5 Format of Size

This specifies the units in which size-related output should be displayed (gigabytes, megabytes, kilobytes, or sectors). The format of size applies to *output only* and is set to megabytes until you reset it.

In properties forms, length values are displayed as a number followed by a `g`, `m`, `k`, or `s` (representing gigabytes, megabytes, kilobytes, or sectors, respectively). If the size cannot be cleanly converted into gigabyte, megabyte, or kilobyte format, it is displayed in sectors instead.

Note – The preferred format of size does not apply to input. Input typically defaults to sectors (`s`), unless gigabytes (`g`), megabytes (`m`), or kilobytes (`k`) are specified. Sizes must be specified as whole numbers.

3.9 Displaying Past and Present Commands

The Command Info Window (FIGURE 3-7) displays commands that are currently being executed by the Visual Administrator, as well as previous commands. Both Volume Manager and system commands are displayed in this window as they are invoked by the Visual Administrator. The status and output of these commands are also displayed. Previously executed commands can be executed again directly from this window.

FIGURE 3-7 Command Info Window

3.9.1 Command Info Sections

The Command Info Window is divided into three sections:

- Command History
- Output of the Highlighted Command
- Commands Running

3.9.1.1 Command History

This section displays a chronological listing (or history) of the commands sent to the Volume Manager or the system for execution. The last ten commands are saved and displayed, with the most recent command at the bottom. If a command is too long, only the first few arguments of the command are shown. Each of these commands is listed along with its status, which is displayed to the right of the command. The status of a command is one of the following:

- `DONE` — Command successfully completed.
- `BROWSE` — Command not executed, just displayed here.
- `ERROR` — Command terminated with error condition.
- `UNKNOWN` — Command status cannot be determined by the Visual Administrator (this rarely occurs, but generally results from an internal Visual Administrator error or a command being interrupted unexpectedly).

3.9.1.2 Output of the Highlighted Command

When a command is highlighted in the Command History section, its output information is displayed in the Output of the Highlighted Command section of the window. The results of the command are indicated here (regardless of whether it succeeded or failed), along with the full command. If the command terminated abnormally or exited with an error condition, the error message is also displayed.

3.9.1.3 Commands Running

This section displays the command that is currently running. This command has been sent to the system or the Volume Manager, but has not yet terminated. As soon as the command completes, it disappears from this section of the window and is placed at the bottom of the Command History section.

3.9.2 Using the Command Info Window

To open the Command Info Window, choose Popup the Command Window from the Options menu. This window may also pop up automatically when an error occurs with a command or operation. To close the window, choose Close from its File menu.

To execute a command directly from the Command Info Window:

1. **Select a command by clicking on the desired command in the Command History window.**

   Once selected, the command is highlighted.

2. **Choose Execute from the Execute menu.**

This procedure sends the selected (highlighted) command to the system or the Volume Manager for execution. This is useful for executing a command again, re-executing a failed command that should now succeed, or executing a command that was only shown (in Browse status) before.

In some circumstances, you can use the Execute with Force item rather than Execute. This option adds `-f` to the executed command, which forces the Volume Manager to complete an operation that is considered unsafe and to disregard error messages. The `-f` option is available with some Volume Manager commands only and does not apply to file system operations.

Caution – The Execute with Force option can result in irreparable loss of data; it should be used only when you are sure that an operation should succeed, even though Volume Manager error checking prevents it.

For detailed information on the Execute option, refer to TABLE 3-10.

3.10 Quitting the Visual Administrator

Most Visual Administrator windows contain a File menu. This menu contains one or both of the following options:

- Close — Close the current window only.
- Exit — Exit both the current window and Visual Administrator completely.

To end a Visual Administrator session, click LEFT on the File option in the menu bar area of the Visual Administrator main window (or any other window) and select Exit. A dialog box appears to confirm that the Visual Administrator session is to be closed completely.

Note – If the window manager is set up with a window menu, accessing that menu from the Visual Administrator main window and selecting Exit will also end the Visual Administrator session. However, closing any other Visual Administrator window through the window manager will only close that particular window.

For detailed information on quitting the Visual Administrator session using the file menu options, refer to TABLE 3-11.

3.11 Quick Reference to Visual Administrator Menus and Forms

The following sections list and describe the menus and forms associated with the topics discussed in this chapter.

3.11.1 Icon Menus

The following menu selections are accessed via the Icon menu.

TABLE 3-3

Menu /Submenu Access	Description
Icon ➤ Maximize Icons	Restores the selected minimized icon(s) to its full size, making it show all of its subicons. **Requirements**: * At least one minimized icon must be selected.
Icon ➤ Minimize Icons	Shrinks the selected icon(s), making it smaller in size. In effect, all of its subicons are hidden. Minimized icons occupy less space and are displayed with their names in reverse type. **Requirements**: * At least one maximized icon must be selected.
Icon ➤ Maximize All Icons	Restores all icons in the current view window, making them show all of their subicons. No icons need to be selected.
Icon ➤ Minimize All Icons	Shrinks all icons in the current view window, making them hide all of their subicons. No icons need to be selected.
Icon ➤ Create Icons	Creates a copy of the icon(s) selected from another view and places the new copy in the current user-created view. Icons that already exist in this user-created view will not be duplicated. **Requirements**: * This option is only available in user-created views.
Icon ➤ Remove Icons	Removes the selected icon(s) from the current user-created view. **Requirements**: * This option is only available in user-created views.
Icon ➤ Help	Accesses a Help window that displays relevant information on icon menu operations.

3.11.2 Projection Menus

Icon projection provides you with visual information about the relationships between icons. When projection is started for an icon, all other icons (representing Volume Manager 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.

Projection operations are accessed through the Projection menu. This menu is located in view windows such as View of `rootdg`. The Projection menu can be used to start or stop projection, as well as to highlight any free subdisk icons.

You can also start or stop Projection by pressing Shift-MIDDLE with the pointer positioned on the desired icon. The following menus are accessed via the Projection menu.

TABLE 3-4

Menu/Submenu Access	Description
Projection ➤ Icon Projection	Provides access to projection options used to start or stop projection for icons.
Projection ➤ Icon Projection ➤ Start	Starts projection for the selected icon(s). When projection is started, all icons related to the selected icon(s) 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. * The selected icon(s) must be associated with at least one other icon in order for projection to take effect.
Projection ➤ Icon Projection ➤ Stop	Stops projection for the selected icon(s). When projection is stopped, all icons related to the selected icon(s) lose their projection highlighting. **Requirements**: * At least one icon must be selected. If the selected icon is not undergoing projection, the Visual Administrator ignores the stop request.
Projection ➤ Icon Projection ➤ Stop All	Stops projection for all icons that are currently undergoing projection.
Projection ➤ Icon Projection ➤ Help	Accesses a Help window that displays relevant information on icon projection operations.

TABLE 3-4 *(Continued)*

Menu/Submenu Access	Description
Projection ➤ Show Free Subdisks	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 Volume Manager 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 Visual Administrator sessions. From the Show Free Subdisks menu, a cascading menu enables you to indicate whether or not to highlight free subdisks.
Projection ➤Show Free Subdisks ➤ Start	Starts highlighting free subdisks immediately and continue to do so until instructed to stop.
Projection ➤Show Free Subdisks ➤ Stop	Stops highlighting free subdisks.
Projection ➤Show Free Subdisks ➤ Help	Accesses a Help window that displays relevant information on show free subdisks operations.
Projection ➤ Help	Accesses a Help window that displays relevant information on projection operations such as disks and show free subdisks.

3.11.3 Analysis Menus

Analysis is the Visual Administrator's way of displaying statistics on the performance of various Volume Manager objects.

Statistics are displayed both visually (using color or pattern) and numerically (using pop-up statistics forms).

Analysis operations are accessed through the Analyze menu. This menu is located in view windows such as View of `rootdg`. The Analyze menu can be used to start or stop analysis, as well as to set analysis-related preferences.

The following menus are accessed through the Analyze menu.

TABLE 3-5

Menu/Submenu Access	Description
Analyze ➤ Start	Begins analysis of the selected icon(s). These icons are added to the list of objects being analyzed. Only volume and VM disk icons can be analyzed. Once analysis is activated, the selected icons begin to display information about their performance characteristics using color or bitmap patterns. **Requirements**: * At least one volume or VM disk icon must be selected.
Analyze ➤ Stop	Terminates analysis of the selected icon(s). These icons are removed from the list of objects being analyzed. When analysis stops, the selected icons return to their pre-analysis states. When analysis is stopped for one icon, other icons undergoing analysis are not affected. **Requirements**: * At least one volume or VM disk icon must be selected. * The selected icon(s) must be undergoing analysis.
Analyze ➤ Analyze All	Begins analysis on all volumes and VM disk icons in all view windows. Once analysis is activated, all icons begin to display information about their performance characteristics using color or bitmap patterns.
Analyze ➤ Stop All	Automatically terminates analysis of all icon(s) in all view windows. All icons return to their pre-analysis states. All colors and patterns disappear. **Requirements**: * Analysis must be in effect.
Analyze ➤ Parameters	Accesses the Analysis Parameters form, which is used to set your preferences for how analysis is to be conducted. **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. **See Also**: * Analysis Parameters Form (TABLE 3-6)
Analyze ➤ Help	Accesses a Help window that displays relevant information on analyze operations.

3.11.4 Analysis Parameters Form

This form, which is accessed via the Analyze menu, is used to set user preferences for conducting analysis.

TABLE 3-6

Field	Description
Sample Rate:	Determines the time interval between data samples. This field is divided into two sections: the slider bar is used to select the interval (1-60) and the menu to its right is used to select 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 in I/Os per second the high and low values that decide the coloring (or pattern) of the volume icons.
Disk Parameters:	Specifies in I/Os per second the high and low values that decide the coloring (or pattern) of the VM disk icons.
Subdisk Parameters:	Specifies in I/Os per second the high and low values that decide the coloring (or pattern) of the subdisk icons.
Log File:	The name of the file to be used for the statistics log. If the file does not already exist, it will be created. The file name is taken to be relative unless a path name is given. To stop logging to the file, the file name text in this field must be erased. This field is optional. This log file is a binary file. To view the log file, `/opt/SUNWvxva/bin/vxvalog2txt` *filename* must be run on this file to process it for viewing.

3.11.5 Analysis Statistics Form

Click on the desired icon that is being analyzed. This form displays analysis statistics relevant to the selected volume or VM disk icon. This form applies only to volume or disk icons that are undergoing analysis. All fields in this form are read-only and cannot be changed.

TABLE 3-7

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.

TABLE 3-7 *(Continued)*

Field	Description
Total Blocks:	The total number of blocks read from or written to the object during the last interval.
Avg Read Time (ms):	The average time, in milliseconds, that it took for a read operation to complete. This is equal to the number of reads during the last interval divided by the total time spent on reads.
Avg Write Time (ms):	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 (secs):	The actual time, in seconds, since the last data was sampled. This may vary slightly from the specified interval time due to uncontrollable variances from system to system.

3.11.6 Options Menus

User preferences for how the Visual Administrator should operate are set via the Options menu. After you have set your preferences, they are saved to the file `$HOME/.vxva_pref`. You should not edit this file. Once set, these preferences are maintained across future Visual Administrator runs.

The following menus are accessed via the Options menu.

TABLE 3-8

Menu /Submenu Access	Description
Options ➤ Show Command	Specifies whether the Command Info Window is to be displayed before each command is executed or only upon command failure.
Options ➤ Show Command ➤ Show on Error	Displays the Command Info Window only when a Volume Manager or system command has failed. This is the default behavior.
Options ➤ Show Command ➤ Show at Start	Displays the Command Info Window whenever a Volume Manager or system command is ready to be sent. Enables you to view the actual command(s) being executed via Visual Administrator.
Options ➤ Show Command ➤ Help	Accesses a Help window that displays relevant information on show command operations.
Options ➤ When Commands Are Ready	Specifies whether commands should be executed immediately or simply displayed for your review.

TABLE 3-8 *(Continued)*

Menu /Submenu Access	Description
Options ➤ When Commands Are Ready ➤ Execute Commands	Automatically executes commands as soon as they are issued. This is the default behavior.
Options ➤ When Commands Are Ready ➤ Show Commands Only	Displays commands in the Command Info Window for your review rather than executing them. Upon approval, you can execute the displayed command directly from the Command Info Window by highlighting that command and then using the Execute menu.
Options ➤ When Commands Are Ready ➤ Help	Accesses a Help window that displays relevant information on command ready operations.
Options ➤ Logging	Starts or stops logging of Visual Administrator commands. Logging records all commands sent to the Volume Manager or the system by the Visual Administrator in a specified file. From the Logging menu, a cascading menu enables you to indicate whether logging should be activated or deactivated. Unlike other user preferences, the logging setting is not saved across Visual Administrator sessions. **See Also**: * Log File Form (TABLE 3-9)
Options ➤ Logging ➤ Start	Begins recording all commands to a log file. A log file is created if one does not already exist. The file to be used for logging must be specified in the resulting form. If the file exists, you must have permission to write to that file. The log information will be appended to the end of the specified file.
Options ➤ Logging ➤ Stop	Stops recording all commands to the log file. When logging is discontinued, you are responsible to remember the name of the log file that was used.
Options ➤ Logging ➤ Help	Accesses a Help window that displays relevant information on logging operations.
Options ➤ Popup the Command Window	Accesses and displays the Command Info Window. This window displays current and previous commands, along with the status of each command. Once accessed in this way, the Command Info Window remains visible until it is closed via its File menu. See "Displaying Past and Present Commands" for further details.

TABLE 3-8 *(Continued)*

Menu /Submenu Access	Description
Options ➤ Format of Size	Specifies the units (gigabytes, megabytes, kilobytes, or sectors) to be used for size-related output. The unit of size is set to megabytes until you reset it. In properties forms, length values are displayed as gigabytes (g), megabytes (m), kilobytes (k), or sectors (s). If the size cannot be cleanly converted into gigabytes, megabytes, or kilobytes, it is displayed in sectors instead (even though another format of size preference may be set). The preferred format of size applies to output only and does not impact input in any way. Input typically defaults to sectors, unless gigabytes, megabytes, or kilobytes are specified.
Options ➤ Format of Size ➤ Gbytes	Use gigabytes when displaying size-related output.
Options ➤ Format of Size ➤ Mbytes	Use megabytes when displaying size-related output.
Options ➤ Format of Size ➤ Kbytes	Use kilobytes when displaying size-related output.
Options ➤ Format of Size ➤ Sectors	Use sectors when displaying size-related output.

3.11.7 Options Log File Form

The Log File Form accessed via the Options menu, is used for logging purposes.

TABLE 3-9

Field	Description
Log File:	The name of the file (and path name) to be used to store the command log. If no path is specified, the file is created in the directory from which the Visual Administrator session was started. **Requirements**: * You must have permissions appropriate to access and write to the named file (and any directories in its path). * If a path name is included, it must be valid.

3.11.8 Execute Menus

In addition to viewing or previewing commands, the Command Info Window can be used to execute commands directly through its Execute menu. This is useful for re-executing commands or for executing commands that were previously only shown in this window.

Caution – The **Execute With Force Option** can cause irreparable loss of data and should only be used when absolutely necessary.

TABLE 3-10

Menu /Submenu Access	Description
Execute ➤ Execute	Execute the command highlighted in the Command History section of the window. Visual Administrator sends the highlighted command to the Volume Manager or the system for execution. **Requirements**: * A single command in the Command History section of the window must be highlighted.
Execute ➤ Execute With Force Option	**Caution - This operation can cause irreparable loss of data and should only be used when absolutely necessary.** Forcefully executes the Volume Manager command highlighted in the Command History section of the window (using the `-f` option). This option effectively forces the Volume Manager to complete an operation that is considered unsafe and should therefore be used only when you are certain that an operation should be performed in this way. **Requirements**: * The force option works for some Volume Manager commands only (those that take the `-f` option). The `-f` option does not apply to file system operations offered through Visual Administrator. * This option should only be used when you are sure that an operation should succeed, even though Volume Manager error checking prevents it.

3.11.9 File Menus

Most Visual Administrator windows contain a File menu, which is used to exit that particular Visual Administrator window and/or to exit the Visual Administrator session completely.

The following menu selections are accessed via the File menu.

TABLE 3-11

Menu /Submenu Access	Description
File ➤ Close	Closes the current window only.
File ➤ Exit	Both closes the current window and exits Visual Administrator completely. Although the Visual Administrator session is closed down, user preferences set during the session are retained. Since Visual Administrator operations are applied to the Volume Manager configuration as they are issued, quitting a Visual Administrator session has no effect on the configuration and does not undo any changes made through Visual Administrator.

CHAPTER 4

Disk Operations

4.1 Introduction

Disk operations fall into the following basic categories:

- *Physical disk* — a physical disk drive that may or may not be under Volume Manager control. Physical disks appear as cylindrical icons labeled `PD`.
- *VM disk* — a disk that has been both placed under Volume Manager control and added to a disk group. VM disks appear as cylindrical icons labeled `D`.
- *Disk group* — a group of VM disks that share a common configuration. Disk groups are represented by "View of *disk_group*" windows.

Refer to Chapter 1, "Introduction to the Volume Manager" for a more detailed explanation of each category.

The quick reference sections at the end of this chapter provide information on Visual Administrator disk and disk groups menus and forms.

4.1.1 Disk Naming

When performing disk administration, you should understand the difference between a *device name* and a *disk name.*

The *device name* (sometimes referred to as *devname* or *disk access name*) is the location of the disk. The syntax of a device name is `c#t#d#s#`, where:

- `c#` is the number of the controller to which the disk drive is attached.
- `t#` is the number of the target disk on that controller.
- `d#` is the number of the disk (or UNIX partition).

- `s#` is the number of the disk slice.

The full path name of a device is `/dev/vx/dmp/`*devicename* . In this document, only the device name is listed and `/dev/vx/dmp` is assumed. An example of a device name is `c0t0d0s2`.

The *disk name* (sometimes referred to as *disk media name*) is an administrative name for the disk, such as `disk01`. If you do not assign a disk name, the disk name defaults to `disk`## if the disk is being added to `rootdg` (where ## is a sequence number). Otherwise, the default disk name is *groupname##,* where *groupname* is the name of the disk group to which the disk is added.

4.1.2 Accessing Disk Operations

The Visual Administrator offers two levels of disk and disk group operations:

- Basic disk operations available through the Basic-Ops menu.
- Advanced disk and disk group operations available through the Advanced-Ops menu.

Note – Some disk operations are not available in demo mode.

4.2 Physical Disk Operations

Physical disks that have been initialized for use by the Volume Manager appear in the View of Disks window. FIGURE 4-1 shows a View of Disks window containing several physical disks. Some of the physical disk icons contain shaded partitions. The VM disk associated with the shaded partition is part of a disk group.

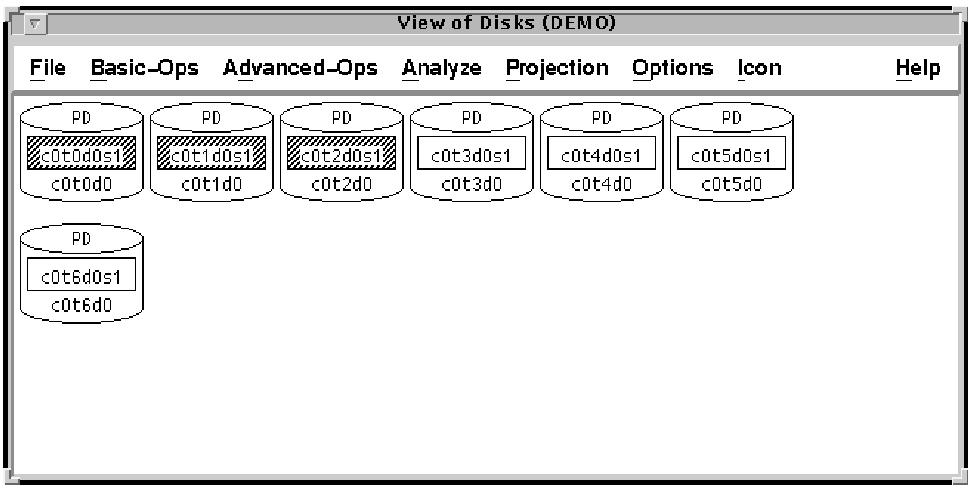

FIGURE 4-1 View of Disks

4.2.1 Identifying a Physical Disk to the Volume Manager

Before a disk can be used within the Volume Manager, it must be identified to the Volume Manager. To be used as a VM disk, it must also be added to a disk group.

The Visual Administrator enables you to identify a disk to the Volume Manager in two ways:

- Using the basic approach, which also partitions the disk and adds it to a disk group.
- Using the advanced approach, which does not automatically partition the disk or add it to a disk group.

4.2.1.1 Basic Approach

The following underlying steps are handled automatically when a disk is placed under Volume Manager control via the Visual Administrator's Basic-Ops menu:

- Initializing, and partitioning the raw disk
- Initializing the new disk for Volume Manager use (installing header and configuration information on the disk)
- Adding the disk to a disk group

Caution – This operation alters the partitioning of a disk and should therefore be used with caution.

To make a physical disk available to the Volume Manager:

1. **Go to the View of Disks window.**
2. **From the Basic-Ops menu, choose Disk Operations, then Add Disks.**

   The Add Disks form is displayed.
3. **Complete the Add Disks Form.**

   For detailed information on completing this form, refer to online help or see TABLE 4-2.
4. **When the form is properly completed, click Apply to activate the disk add operation.**

   After a confirmation window pops up, a new physical disk icon (with the name supplied as the new disk name) is displayed. The new physical disk icon contains a partition icon. The partition icons of physical disks under Volume Manager control are patterned or colored to differentiate them from others. If a disk group was specified, a new VM disk icon is displayed in the view window corresponding to that disk group.

4.2.1.2 Advanced Approach

The Advanced-Ops menu provides an option to initialize a new disk for Volume Manager use (installing header and configuration information on the disk). If the disk is to be partitioned and/or added to a disk group, these tasks must be performed separately afterwards.

To make a physical disk available to the Volume Manager:

1. **Go to the View of Disks window.**
2. **From the Advanced-Ops menu, choose Disk, then Initialize.**

   The Disk Init form is displayed. This form is used to initialize regions of a disk used by the Volume Manager.
3. **Complete the Disk Init Form.**

   For detailed information on completing this form, refer to online Help or see TABLE 4-5.

4. **When the form is properly completed, click Apply to activate the disk initialization.**

   A new physical disk icon containing a partition icon is displayed.

4.2.2 Adding a Disk to a Disk Group

To add a disk to the `rootdg` disk group:

1. **Access the View of `rootdg` window and position it so that it is at least partially visible.**
2. **Go to the View of Disks window and drag the partition icon corresponding to the physical disk to be added to a disk group into the View of `rootdg` window. Drop the partition icon anywhere within the View of `rootdg` window.**

   Do not attempt to drag the entire physical disk icon surrounding the partition.

   A new VM disk icon is displayed in the View of `rootdg` window. The partition icon in the View of Disks window changes color or pattern to indicate that it now belongs to a disk group.

   This procedure can be used to add a physical disk to any other disk group by substituting the appropriate disk group view for View of `rootdg`.

4.2.3 Replacing a Disk

Disks most often need replacing when they fail or start to behave strangely. Replacing an existing VM disk with a new one (if the disk has failed, for instance) involves a set of procedures including the following:

- Initializing and partitioning the raw disk
- Initializing the new disk for Volume Manager use (installing header and configuration information on the disk)
- Replacing the old disk with the new one (connecting the existing disk media record to the new disk)

These underlying steps are handled automatically when the disk replacement is performed through the Visual Administrator's Basic-Ops menu.

Caution – This operation alters the partitioning of a disk and should therefore be used with caution.

The disk replacement procedure outlined here can only be performed on a disk that has failed. The Visual Administrator identifies such a disk by altering the disk icon's color or pattern when failure occurs.

To replace a VM disk:

1. **Go to the View window corresponding to the disk group containing the VM disk to be replaced and select the failed disk by clicking its icon.**

   The Visual Administrator should portray the failed disk icon in a different color or pattern.

2. **From the Basic-Ops menu, choose Disk Operations, then Replace Disks.**

   The Replace Disks form is displayed.

3. **Complete the Replace Disks Form.**

   For detailed information on completing this form, refer to online help or see TABLE 4-3.

4. **When the form is properly completed, click Apply to activate the disk replacement.**

   If the replacement is successful, the VM disk icon associated with the failed physical disk returns to a normal state (the special icon coloring or pattern goes away). In the View of Disks window, a physical disk icon with a new name and a new partition is displayed.

4.2.4 Moving Data From a Failed Physical Disk to a Different Physical Disk

If a physical disk fails or starts to fail, you can move the data from the failed physical disk to a different physical disk as follows:

1. **Determine which physical disk has failed.**

   You may be able to determine which physical disk has failed by clicking the button over the failed VM disk icon to access the VM Disk Properties form. The Visual Administrator portrays a failed VM disk icon in a different color or pattern.

   The `Physical Disk` text field on the VM Disk Properties form gives the device name for the failed disk that the VM disk is assigned to. Write down the device name of the failed disk; you will need it later to locate the failed physical disk so that you can replace it. For detailed information on this form, refer to online help or see TABLE 4-11.

2. **Evacuate subdisks from the failed disk as follows:**

Note – This operation can only be performed between two disks in the same disk group.

---

   a. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

   b. **Select the disk *from* which subdisks are to be evacuated by clicking on its icon.**

   c. **From the Basic-Ops menu, choose Disk Operations, then Evacuate Subdisks.**

      The Evacuate Subdisks form is displayed.

   d. **Complete the Evacuate Subdisks Form.**

      For detailed information on completing this form, refer to online Help or see TABLE 4-4.

   e. **Click Apply to activate the subdisk evacuation.**

      The subdisks move from their original disk icon to the targeted disk icon.

3. **Refer to the device name that you wrote down earlier in this procedure to find the physical disk that needs to be replaced.**

4.2.5 Displaying Properties for a Physical Disk

The properties of the physical disk corresponding to a disk icon can be displayed in a properties form. Properties forms provide detailed information about the characteristics of a particular disk.

To access the properties form for a particular physical disk or partition icon:

- **Click on its icon. (If the VM disk is under analysis, press Shift-RIGHT instead.)**

The appropriate properties form displays details of the physical disk or partition configuration. Physical disk and partition properties are read-only and cannot be changed through their properties forms. For detailed information on the physical disk properties form, refer to online Help or TABLE 4-12. For detailed information on the partition properties form, refer to online Help or see TABLE 4-13.

4.3 VM Disk Operations

VM disks are accessible through disk group view windows.

4.3.1 Determining Amount of Free Space on VM Disks

For operations that require free disk space, the Volume Manager and Visual Administrator usually enable you to designate disks with sufficient free space. If no disks are specified, disks with available space are automatically used if an operation is performed with Basic-Ops. With Advanced-Ops, however, you are expected to designate disks with sufficient space. Before designating a disk to be used, verify the amount of free space on that disk.

To check the available space on a disk:

- **Click on a portion of the Volume Manager disk icon that is not covered by another icon.**

The VM disk properties form is displayed. At the bottom of this form is a field that indicates the maximum free space available for use on this disk. For detailed information on this form, refer to TABLE 4-11.

It is sometimes possible to obtain information about a specific region of a VM disk that contains free space. If subdisks have been removed from a disk, resulting in visual gaps (or “holes”) between the subdisk icons located on that disk, you can check the amount of space freed by the subdisk removal.

FIGURE 4-2 illustrates a VM disk icon containing a hole created by the removal of a subdisk.

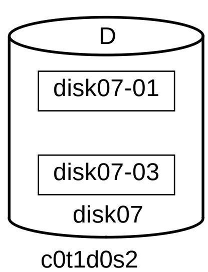

FIGURE 4-2 VM Disk with a Hole

To check the amount of free space in a specific region (hole) of a disk:

1. **Go to the view window corresponding to the disk group in which you want to perform the operation (View of `rootdg`, by default).**
2. **Click the RIGHT button on the gap between subdisk icons on the desired VM disk icon.**

   The Free Space form is displayed. This form provides read-only information about the amount of free space in the hole. For detailed information on this form, refer to online help or see TABLE 4-8.
3. **Click on Cancel to close the Free Space form.**

4.3.2 Removing a VM Disk From a Disk Group

A VM disk can be removed from a disk group when it is no longer necessary to access the physical disk that the VM disk is assigned to or when that physical disk is about to be removed from Volume Manager control altogether. A physical disk cannot be removed until its VM disk has been removed from its disk group. Any subdisk residing on the VM disk that is assigned to that physical disk must be removed or moved to another VM disk before the VM disk can be removed from its disk group.

Note – The last VM disk in a disk group cannot be removed. In order for the last VM disk in a disk group to be removed, the disk group itself must be deported as described in Section 4.4.2, "Deporting a Disk Group." Note, however, that the last VM disk can never be removed from the `rootdg` disk group, nor can the `rootdg` disk group be deported.

To remove a VM disk from the `rootdg` disk group:

1. **Go to the View of `rootdg`.**
2. **Select the VM disk to be removed by clicking on its icon.**
3. **From the Advanced-Ops menu, choose Disk Group, then Remove Disks.**

   The VM disk icon is removed from the View of `rootdg`, indicating that the VM disk no longer belongs to this disk group.

   This procedure can be used to remove a disk from any other disk group by substituting the appropriate disk group view for View of `rootdg`.

4.3.3 Adding a VM Disk to the Hot-Relocation Pool

Hot-relocation is the ability of a system to automatically react to I/O failure by relocating redundant subdisks to other disks and restoring the affected VxVM objects and data. If a disk has already been designated as a spare in the disk group, the subdisks from the failed disk are relocated to the spare disk. Otherwise, any suitable free space in the disk group is used. See Chapter 1, “Introduction to the Volume Manager” for more information.

To designate a disk as a spare:

1. **Go to the disk group view containing the VM disk to be designated as a spare.**
2. **Open the properties form for the VM disk by clicking the RIGHT button on its icon.**

   For detailed information on this form, refer to TABLE 4-11.

3. **Set the** `Spare` **field in the VM disk’s properties form to** `Yes`**.**
4. **Click Apply to activate the spare operation.**

   The VM disk icon’s label should become `S`. Any VM disk in this disk group can now use this disk as a spare in the event of a failure.

   If a disk fails, the failing disk’s VM disk icon should change color/bitmap pattern and hot-relocation should automatically occur (if possible). You should be notified of the failure and relocation via electronic mail. After successful relocation, you may want to replace the failed disk (as described in Section 4.2.3, “Replacing a Disk”).

4.3.4 Removing a VM Disk From the Hot-Relocation Pool

While a disk is designated as a spare, the space on that disk is not used as free space for the creation of VxVM objects within its disk group. If necessary, you can free a spare disk for general use by removing it from the pool of hot-relocation disks.

To remove a disk from the hot-relocation pool, follow the procedure for adding a VM disk to the hot-relocation pool, but change the `Spare` field to `No.`
The VM disk icon’s label should become `D` again.

4.3.5 Displaying and Altering Properties for a VM Disk

The properties of the VM disk corresponding to a disk icon can be displayed (and potentially adjusted) in a properties form. Properties forms provide detailed information about the characteristics of a particular disk.

To access the properties form for a particular VM disk:

- **Move the pointer onto that icon and click the RIGHT button. (If the VM disk is under analysis, press Shift-RIGHT instead.)**

  The appropriate properties form displays details of the VM disk configuration. Some of the VM disk's properties can be altered directly through its properties form by altering the appropriate fields and then selecting the form's Apply button. For detailed information on this form, refer to online help or see TABLE 4-11.

4.4 Disk Group Operations

Disk groups are represented by the Visual Administrator as view windows (such as `rootdg`) rather than icons.

4.4.1 Initializing a Disk Group

A disk group must contain at least one disk at the time it is created.

To create a new disk group:

1. **Go to the View of Disks window.**
2. **Select the partition icon corresponding to the physical disk to be added to the new disk group.**
3. **From the Advanced-Ops menu, choose Disk Group, then Initialize.**

   The Initialize Disk Group form is displayed.
4. **Complete the Initialize Disk Group form, entering a unique name to be applied to the new disk group.**

   For detailed information on this form, refer to online help or see TABLE 4-15.
5. **Click Apply to activate the disk group initialization.**

A button representing the newly-created disk group is displayed in the Visual Administrator main window. You should now be able to access the new disk group's view by clicking on its button.

4.4.2 Deporting a Disk Group

Deporting a disk group disables access to that disk group. A disk group may be deported so that its last VM disk can be removed, thus enabling the physical disk assigned to that VM disk to be reused.

Note – The default disk group, `rootdg`, cannot be deported.

Once deported, the disk group is inaccessible and any VM disk icons that belonged to that disk group are dismissed along with the disk group view. The partition icon corresponding to a VM disk that was dismissed when its disk group was deported reverts to an unshaded state. However, the partition retains its disk access record and knowledge of the deported disk group until it is reused (assigned to another disk group) or removed.

To deport a disk group:

1. **Go to the view of the disk group to be deported.**
2. **From the Advanced-Ops menu, choose Disk Group, then Deport Disk Group.**

   The Deport Disk Group form is displayed. For detailed information on this form, refer to online help or see TABLE 4-17.
3. **Ensure that the Deport Disk Group form displays the current disk group name as the disk group to be deported, then click Apply.**

   The disk group view is dismissed, along with any VM disks that it contained. The view button corresponding to the deported disk group is also dismissed from the Visual Administrator main window.

4.4.3 Importing a Disk Group

Importing a disk group enables access to a disk group that has been deported. To import a deported disk group, its former name must be known and remain unused. In addition, at least one partition (containing a disk access record) formerly assigned to the deported disk group must remain unused. If all disks associated with a deported disk group have since been reused, that disk group cannot be imported.

To import a disk group:

1. **Go to any view window.**
2. **From the Advanced-Ops menu, choose Disk Group, then Import Disk Group.**

   The Import Disk Group form is displayed.

3. **Complete the Import Disk Group form, entering the name of the disk group to be imported.**

   The named disk group must have been deported at one time. At least one partition formerly assigned to the deported disk group must still exist (unused) in order for the disk group to be importable. If you want to rename the disk group on import, you must complete the New name field. For detailed information on this form, refer to online help or see TABLE 4-16.

   Indicate whether you want to clear the host ID.

   Note – This option clears the existing host ID that is recorded on all disks in the disk group.

   Use this option only in one of the following two cases:

   - After you have physically moved disks between machines and are in the process of importing the disk group onto a new host.
   - If you are certain that only one machine (though two machines are dual porting the same disk) is attempting to access a particular disk group at a given time.

4. **Click Apply to activate the disk import.**

   The view button corresponding to the disk group that has been imported is displayed in the Visual Administrator main window. The partition icon(s) corresponding to the VM disk icon(s) that formerly belonged to that disk group and has not been reused becomes shaded. When the view of the newly-imported disk group is accessed, its VM disk icon(s) is visible.

4.5 DMP Operations

This is a set of utilities for administering the DMP driver.

4.5.1 Getting the DMP Node

This is a utility which gets the DMP node given a disk name from `/dev/rdsk`. It also shows all the paths to the DMP node and their current status.

To get a DMP node, do as follows:

1. **Go to any view window.**

2. **From the Advanced-Ops menu, select DMP, then DMP node.**

   The Disk Name Entry form appears.

3. **Enter the disk name and select apply.**

   For detailed information on this form, refer to TABLE 4-9.

4. **The Disk Name Form disappears and a new form comes up.**

   This form shows the DMP node and the list of paths available to it with the status of each path. If the DMP node is not under Volume Manager control, the lists of paths available is not shown. For detailed information on this form refer to TABLE 4-10.

4.5.2 Reconfiguring DMP

This is another utility which enables the user to reconfigure the DMP database.

To reconfigure the DMP database, do as follows:

1. **Go to any view window.**

2. **From the Advanced-Ops menu, select DMP, then Reconfigure DMP.**

   This reconfigures the DMP database.

4.6 Quick Reference to Disk Menus and Forms

This section provides information on menus and forms relating to disk operations. The Basic-Ops and Advanced-Ops menus provide access to disk-related operations. Most of the disk-related menus are accessed through the Advanced-Ops menu.

Most menus provide a Help selection, which contains information relevant to the items and operations listed in that particular menu.

Note – Some disk operations are not available through demo mode.

4.6.1 Disk Menus

The following menus and menu selections are accessed via the Basic-Ops or Advanced-Ops menus:

TABLE 4-1

Menu/Submenu Access	Description
Basic-Ops ➤ Disk Operations	Provides a cascading menu to add, replace, and evacuate disks.
Basic-Ops ➤ Disk Operations ➤ Add Disks	Adds a disk to the Volume Manager, placing it under Volume Manager control. This involves initializing and partitioning the raw disk; initializing the disk for Volume Manager use; and adding the disk to a disk group. **See Also**: * Add Disks Form (TABLE 4-2).
Basic-Ops ➤ Disk Operations ➤ Replace Disks	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 Volume Manager 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. **See Also**: * Replace Disks Form (TABLE 4-3).
Basic-Ops ➤ Disk Operations ➤ Evacuate Subdisks	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. **See Also**: * Evacuate Subdisks Form (TABLE 4-4).
Basic-Ops ➤ Disk Operations ➤ Help	Accesses a Help window that displays relevant information on basic disk operations.
Advanced-Ops ➤ Disk	Provides a cascading menu to initialize, define, online, and offline disks.

TABLE 4-1 *(Continued)*

Menu/Submenu Access	Description
Advanced-Ops ➤ Disk ➤ Initialize	Identifies a disk to the Volume Manager and initializes the disk for Volume Manager 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. **Requirements**: * The disk should not already be initialized. **See Also**: * Disk Init Form (TABLE 4-5).
Advanced-Ops ➤ Disk ➤ Define	Defines a disk access record, which enables the Volume Manager to scan the disk. This makes the disk accessible, but does not initialize the disk. **See Also**: * Define Disk Form (TABLE 4-6).
Advanced-Ops ➤ Disk ➤ Remove	Removes the VM disk associated with the selected partition(s) from Volume Manager 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 Volume Manager control. **Requirements**: * At least one partition icon corresponding to a VM disk must be selected. * The VM disks corresponding to the selected partition(s) cannot belong to a disk group at the time of removal.
Advanced-Ops ➤ Disk ➤ Online	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. **See Also**: * Disk Online Form (TABLE 4-7).
Advanced-Ops ➤ Disk ➤ Offline	Places the disk access record on the selected partition(s) 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 a VM disk).
Advanced-Ops ➤ Disk ➤ Help	Accesses a Help window that displays relevant information on advanced disk operations.
Advanced-Ops ➤ DMP ➤ Get DMP Node	Given a disk name from `/dev/rdsk`, this option gives the DMP node, list of paths, and their status.
Advanced-Ops ➤ DMP ➤ Reconfigure DMP	Reconfigures the DMP database.

4.6.2 Disk Forms

Some disk operations result in the appearance of forms, which must be completed in order for that operation to proceed. Most forms provide online help, which contains information relevant to the fields and other aspects of that form.

Some form fields are required while others are not. Fields in these forms are read / write fields, unless listed as read only.

4.6.2.1 Add Disks Form

This form is used to place a disk under Volume Manager control.

TABLE 4-2

Field	Description
New disk name:	Specify the name of the new *physical* disk in `c#t#d#` form. The name must be unique within this disk group.
Disk group:	Type 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.

4.6.2.2 Replace Disk Form

This form is used to replace an existing VM disk that has failed.

TABLE 4-3

Field	Description
Old VM disk name:	Displays the name of the failed VM disk in this disk group. It is read only and cannot be changed.
New physical disk name:	Type the name of the new physical disk that is to replace the existing one. The name should be in `c#t#d#` form and must be unique.

4.6.2.3 Evacuate Disk Form

This form is used to transfer subdisks from one VM disk to another.

TABLE 4-4

Field	Description
Disk group name:	Type the name of the disk group to which both disks belong. Both disks must share the same disk group.
Evacuate From:	Displays the name of the VM disk from which the subdisks are to be evacuated, which can be changed. **Requirements**: * Both disks must belong to the same disk group. * Subdisks must exist on the disk from which the evacuation is to be conducted.
To:	Type the name of the VM 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).

4.6.2.4 Disk Init Form

This form is used to initialize a disk for Volume Manager use.

TABLE 4-5

Field	Description
Public device:	Type the name of the device node that represents a partition available for use. This name must be a valid entry in `/dev/vx/rdmp/`. A name of the form `c#t#d#s#` is appropriate.
Device type:	Select the desired disk type. The `simple` type 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 (default) assumes that the public and private regions are stored on different disk partitions. The `nopriv` type has no private region; thus log and configuration copies cannot be written to the disk.
Public Length (0 for whole device):	Indicate the length of the public section of the disk. If zero is provided as the length, the Volume Manager computes a default value from available partition table information. This length must be valid and cannot exceed the length of the disk.

TABLE 4-5 *(Continued)*

Field	Description
Private Length (0 for default):	Indicate the length of the private region of the disk. When one is not specified, the Volume Manager chooses a default value. This length must be valid and cannot exceed the length of the disk. This field is optional.
Number of config copies:	Indicate the number of configuration copies to be stored in the private section of this disk. The default value is 1 copy.
Comment:	Type a comment appropriate for this Volume Manager disk. The maximum length of the comment is 40 characters. This field is optional.

4.6.2.5 Disk Define Form

This form is used to define a disk.

TABLE 4-6

Field	Description
Public device:	Type the path name of the device node that represents a partition available for use. This name must be a valid entry in `/dev/vx/rdmp/`. A name in the form `c#t#d#s#` is appropriate.
Device type:	Select the desired disk type. The `simple` type 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 (default) 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):	Indicate the length of the public section of the disk. If zero is provided as the length, the Volume Manager computes a default value from available partition table information. This length must be valid and cannot exceed the length of the disk.
Offline:	Indicate whether to initially place the disk in the offline state. The default is `No`.
Comment:	Type a comment appropriate for this Volume Manager disk. The maximum length of the comment is 40 characters. This field is optional.

4.6.2.6 Disk Online Form

This form is used to place a disk online.

TABLE 4-7

Field	Description
Device name:	Indicate the disk access name of the disk to be placed online. This must be a valid disk access name. This field is required.

4.6.2.7 Free Space Form

This form provides information about a specific region of a Volume Manager disk that contains free space. All fields in this form are read only and cannot be changed. 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 a VM disk icon.

TABLE 4-8

Field	Description
Device:	Displays the name of the VM disk where this free space resides.
Hole offset:	Indicates the offset into the Volume Manager disk where this free space extent begins.
Hole size:	Indicates the size of this free space extent. The units used are specified by the user under the Options menu.

4.6.3 DMP Forms

4.6.3.1 Disk Name Entry Form

TABLE 4-9

Field	Description
Disk Name:	Enter the disk name from `/dev/rdsk`.

4.6.3.2 Path Form

TABLE 4-10

Field	Description
DMP mode name:	Specifies the DMP node.
Paths to Disk:	Lists the paths to the disk and their status.

4.6.4 Properties Forms

Properties forms exist for VM disks, physical disks, and partitions.

4.6.4.1 VM Disk Properties Form

This form provides detailed information on the attributes of a particular VM disk that is under Volume Manager control. The information displayed in this form actually corresponds to the disk media record associated with a disk.

Note – Most fields in this form are read only. Some 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 the Apply button.

TABLE 4-11

Field	Description
VM disk name:	The name of the Volume Manager disk.
Physical disk:	The name of the physical disk that corresponds to this VM disk, or the last associated physical disk if VM disk has been disconnected. This field is read only.
Disk type:	The disk type with which this VM disk was created. This field is read only.
Public region:	The name of the public region of this disk. This field is read only.
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.
Public region offset:	The offset, in sectors, of the public region on the disk. This field is read only.
Private region offset:	The offset, 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.
Public region length:	The length, in sectors, of the public region on the disk. This field is read only.
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.
Disk Attributes:	The attributes of this VM disk. This field is read only.
Spare:	Indicates whether this disk is reserved for use as a hot-relocation spare. This field can be used to designate this disk as a spare (`Yes`) or remove it from the hot-relocation pool (`No`).
Failing:	Indicates whether the disk is currently experiencing some sort of failure. This field is typically set to `No`. When a failure occurs, this changes to `Yes` and the disk cannot be used for space allocation. Although it is possible to force a failing disk to allow space allocation again (by changing this field back to `No`), this is not recommended.
Comment:	The user-specified comment for this VM disk. The maximum length of the comment is 40 characters.
Putil0:	Permanent utility field 0. This is reserved for VM 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.

TABLE 4-11 *(Continued)*

Field	Description
Tutil1:	Temporary utility field 1. This field is reserved, but may be changed.
Tutil2:	Temporary utility field 2. This field is reserved, but may be changed.
Maximum Free Space:	The maximum amount of free space available on this VM disk. This does not take disk extents into account. This number assumes every free sector on the VM disk is usable. In the case of a failing disk, this field will indicate that the available free space is 0. This field is read only.

4.6.4.2 Physical Disk Properties Form

This form provides detailed information on the attributes of a particular physical disk. All fields in this form are read only and cannot be changed.

TABLE 4-12

Field	Description
Device:	The raw device node for this physical disk.
Device type:	A brief description of the device type. Possible device types include SCSI hard drive and Floppy. This field displays the Visual Administrator demo drive when the Visual Administrator is running in demo mode (in which case this physical disk does not correspond to a real disk).
Disk heads:	The number of read/write heads on this disk.
Cylinders:	The number of cylinders on this disk.
Sectors:	The total number of sectors on this disk.
Sector size (bytes):	The size, in bytes, of each sector on this disk.
Total size (sector):	The total size of the disk, in sectors.

4.6.4.3 Partition Properties Form

This form provides detailed information on the attributes of a particular partition. All fields in this form are read only and cannot be changed.

TABLE 4-13

Field	Description
Device:	The device node that the 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, represented in units.
Type:	The identification tag associated with this partition. (The type DEMO represents a partition that does not correspond to a real partition on the system, but is instead just displayed in the Visual Administrator for demo mode purposes.)
Permissions:	Flags that represent the permissions and status conditions of the specified 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.

4.7 Quick Reference to Disk Group Menus and Forms

This section provides information on menus and forms relating to disk group operations. Most disk group operations are carried out in the View of `rootdg` (or some other disk group view) or View of Disks.

The Advanced-Ops menu provides access to disk group related operations. Most menus provide a Help selection, which contains information relevant to the items and operations listed in that particular menu.

4.7.1 Disk Group Menus

The following menus and menu selections are accessed via the Advanced-Ops menu:

TABLE 4-14

Menu /Submenu Access	Description
Advanced-Ops ➤ Disk Group	Provides access to disk group operations.
Advanced-Ops ➤ Disk Group ➤ Initialize	Defines a new disk group with a name you specify. The new disk group contains one or more VM disks corresponding to the partition(s) you selected. **Requirements**: * At least one partition icon must be selected. **See Also**: * Initialize Disk Group Form (TABLE 4-15)
Advanced-Ops ➤ Disk Group ➤ Import Disk Group	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. This option also gives you the choice to clear the existing host ID that is recorded on all disks in the disk group. **See Also**: * Import Disk Group Form (TABLE 4-16)
Advanced-Ops ➤ Disk Group ➤ Deport Disk Group	Disables access to a disk group. A deported disk group is no longer accessible and its view window is dismissed. 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. * The root disk group, `rootdg`, cannot be deported. **See Also**: * Deport Disk Group Form (TABLE 4-17)
Advanced-Ops ➤ Disk Group➤ Add Disk	Adds a VM 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. **See Also**: * Add Disk Form (TABLE 4-18)

TABLE 4-14 *(Continued)*

Menu /Submenu Access	Description
Advanced-Ops ➤ Disk Group ➤ Disconnect Disks	Disables the selected VM 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 VM disk icon must be selected. * The VM disk icon(s) must contain a disk media record at the time of selection.
Advanced-Ops ➤ Disk Group ➤ Reconnect Disks	Enables a VM disk that has previously been disconnected. This involves connecting the selected VM disk's disk media record with the selected disk access record. Although the VM disk must be disconnected, it does not necessarily have to be reconnected to its former partition (disk access record). **Requirements**: * One VM disk icon and one partition icon must be selected. * Neither the VM disk icon nor the partition icon can already be connected.
Advanced-Ops ➤ Disk Group➤ Remove Disks	Removes the selected VM disk(s) from a disk group. Disks are removed from the disk group in which they reside. Any subdisks that exist on the selected disk(s) must be removed before the disk can be removed. **Requirements**: * At least one VM 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. * The root disk group, `rootdg`, cannot be emptied of all disks.
Advanced-Ops ➤ Disk Group ➤ Help	Accesses a Help window that displays information relevant to the disk group operations

4.7.2 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 online help that contains information relevant to the fields and other aspects of that particular form.

Some form fields are required while others are not. Fields in these forms are read / write fields, unless listed as read only.

4.7.2.1 Initialize Disk Group Form

This form is used to name a new disk group consisting of selected disks.

TABLE 4-15

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.

4.7.2.2 Import Disk Group Form

This form is used to make the specified disk group available to the system.

TABLE 4-16

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.
New name:	The disk group will be renamed to the contents of this field when imported. This field does not contain a default value. If you do not fill in this field, the disk group will not be renamed.
Clear host ID:	Use this field with caution to clear the existing host ID that is recorded on all disks in the disk group. The default is `NO`. For details on when this field should be switched to YES, refer to Section 4.4.3, "Importing a Disk Group."

4.7.2.3 Deport Disk Group Form

This form is used to make the specified disk group inaccessible to the system.

TABLE 4-17

Field Name	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.
New name:	This will become the disk group's new name after it is deported. This field does not contain a default value. If you do not fill in this field, the existing name remain the same
New host:	This will set the disk group's host as part of the deport operation. This field does not contain a default value. If you do not fill in this field, the host name will remain unchanged. This field, along with the New name field, are useful for fixing a `rootdg` on another machine that shares disks with the current host.

4.7.2.4 Add Disk Form

This form is used to add a VM disk to a disk group.

TABLE 4-18

Field Name	Description
Disk Group:	The name of the disk group to which the VM disk is to be added. This must be a valid disk group. This field is required.
Disk Media Name:	The name of the VM 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.

CHAPTER 5

File System Operations

The Visual Administrator offers a limited set of file system operations. This chapter provides instructions on performing the following file system operations using the Visual Administrator.

Note – Some operations discussed in this chapter can only be performed with the VERITAS File System (VxFS).

The quick reference sections at the end of this chapter provide information on the Visual Administrator menus and forms associated with file system operations.

5.1 File Systems and the Visual Administrator

The Visual Administrator performs file system operations by executing file system commands directly. The file system must be placed on a volume before file system operations can be performed through the Visual Administrator.

Due to the nature of UNIX file systems, the Visual Administrator cannot detect whether an unmounted file system already exists on a given volume. Only mounted file systems can be detected by the Visual Administrator. Therefore, you must be aware of whether an unmounted file system already exists before performing certain file system operations.

The Visual Administrator has no icons for file systems. However, the mount point name is displayed below the volume icon if a mounted file system exists on that volume. If the mount point is very long, only part of it is displayed under the volume icon.

The Visual Administrator adds entries for new file systems to the `/etc/vfstab` file (file system table). However, entries for removed file systems need to be removed from this file manually.

Note – In demo mode, the underlying volumes are artificial; since file system commands require real entities to succeed, file system operations cannot be simulated by the Visual Administrator in demo mode.

5.1.1 File System Type Restrictions

The Visual Administrator supports several file system types, including `vxfs` and `ufs`. However, some operations are not available with all file system types. For detailed information on operations specific to the VxFS file system, refer to the *VERITAS File System (VxFS) System Administrator's Guide*.

The following operations are available with `vxfs` file systems only:

- Defragment
- Snapshot

In addition, the Resize operation is available with `vxfs` and `ufs` only, but is limited to growing a file system for `ufs`.

Note – If `vxfs` is not installed when `vxva` is started, form fields corresponding to options that are specific to `vxfs` file systems are either excluded from forms or inactive.

5.1.2 VxFS Mount Options

When mounting a VxFS file system, options that affect the performance of the file system can be specified. Using these options, the Visual Administrator provides a selection of two default file system configurations. The Visual Administrator automatically specifies these options to the `mount` command. These configurations are designed to correspond to common user environment types, as follows:

- Server — This is the default. No mount options are passed to `mount`, so the default logging mode `log` is in effect. With `log` mode, all structural changes are logged to disk before the system call returns to the application. Even if a system failure occurs, operations will be completed and nothing will be lost. This is a safe mode, but performance is slower.

- Desktop — The `mincache=closesync` option is passed to `mount`. This mode is particularly useful in desktop environments where users are likely to shut off the power on the machine without halting it first. With the `mincache=closesync` mode in effect, only files that are currently being written when the system crashes or is turned off can lose data. In this mode, any unlogged changes to the file are flushed to disk when the file is closed. No logging options are specified, so `log` is in effect.

For more details on the VxFS file system mount options, refer to the *VERITAS File System (VxFS) System Administrator's Guide*.

5.2 Creating a File System on a New Volume

File systems can be created and placed on volumes, one file system per volume. When the Visual Administrator is used to create a file system via the Basic-Ops, File Systems Operations menu, two tasks are actually performed:

- A volume is created on a Volume Manager disk.
- A file system is created on that volume.

The Visual Administrator handles both of these tasks automatically.

A file system can be created on one of three basic volume types:

- Simple — A concatenated volume whose subdisks are arranged both sequentially and contiguously within a plex. Concatenation enables a volume to be created from multiple regions of one or more disks if there is not enough space for an entire volume on a single region of a disk.
- Striped — A volume with data spread fairly evenly across multiple disks. *Stripe units* are relatively small, equally-sized fragments that are allocated alternately and evenly to the subdisks of a single plex. There should be at least two subdisks in a striped plex, each of which should exist on a different disk. Throughput increases with the number of disks across which a plex is striped. Striping helps to balance I/O load in cases where high traffic areas exist on certain subdisks.
- RAID-5 — A volume that uses striping to spread data fairly evenly across multiple disks in an array and enables independent access. It also stripes parity across all the disks in the array, rather than being concentrated on a single parity disk. Each stripe contains a parity stripe units and data stripe units. In comparison to the performance of striped volumes, throughput of RAID-5 volumes decreases since parity information needs to be updated each time data is written. However, in comparison to mirroring, the use of parity reduces the amount of space required.

5.2.1 Creating a File System on a Simple Volume

To create a file system on a simple, concatenated volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Optionally designate a disk to be used for the new volume.**

   Ensure that the disk has sufficient free space to accommodate the desired length of the new volume by clicking the RIGHT button on an unobscured portion of that disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. For detailed information on the VM disk properties form, refer to TABLE 4-11.

   If there is sufficient free space, select the disk by clicking the LEFT button on its icon. If no disks are selected, disks with sufficient free space are automatically used.
3. **From the Basic-Ops menu, select File System Operations, then Create.**

   A submenu listing basic volume types is displayed.
4. **Select Simple.**

   The Simple Volume/FS Create form is displayed (see FIGURE 5-1). This form creates a file system on a concatenated 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 file system on a new volume.

Note – Fields specific to `vxfs` file systems are only displayed if `vxfs` is installed on the system.

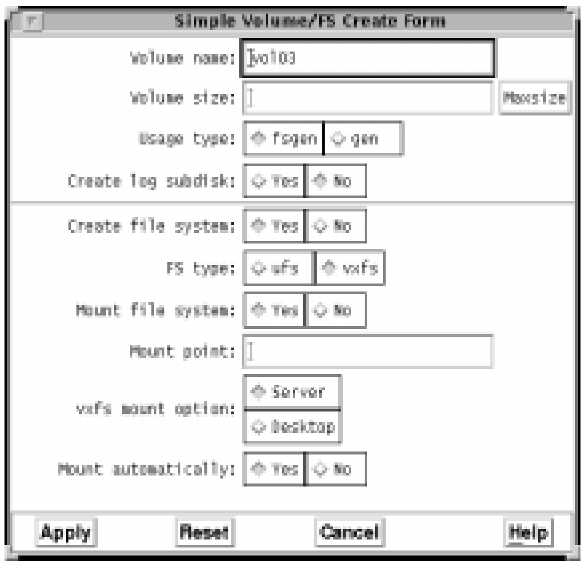

FIGURE 5-1 Simple Volume / FS Create Form

5. **Complete the Simple Volume/FS Create Form.**

   For detailed information on this form, refer to online help or TABLE 5-2.

6. **Click Apply to create the volume and file system.**

   A new volume icon is displayed. Since this is a simple, concatenated volume, it contains a single plex. If the file system is mounted, it is represented by the mount point, which is displayed below the new volume.

   FIGURE 5-2 shows a simple, concatenated volume with a mounted file system.

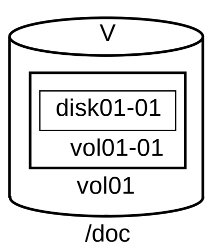

FIGURE 5-2 Simple Volume and File System

5.2.2 Creating a File System on a Striped Volume

To create a file system on a striped volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Optionally, designate at least two disks to be used for the new volume.**

   These disks must belong to the same disk group. Ensure that the disks have sufficient free space to accommodate their portion of the desired length of the new volume by clicking the RIGHT button on an unobscured portion of each disk icon to access its properties form, then checking the value in the Maximum free space field.

   If there is sufficient free space, select the disks by clicking the MIDDLE button on their icons. If no disks are selected, disks with sufficient free space are automatically used.

3. **From the Basic-Ops menu, select File System Operations, then Create.**

   A submenu listing basic volume types is displayed.

4. **Select Striped.**

   The Striped Volume/FS Create form is displayed (see FIGURE 5-3). This form creates a file system on a striped 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 file system on a new volume.

Note – Fields specific to `vxfs` file systems are only displayed if `vxfs` is installed on the system.

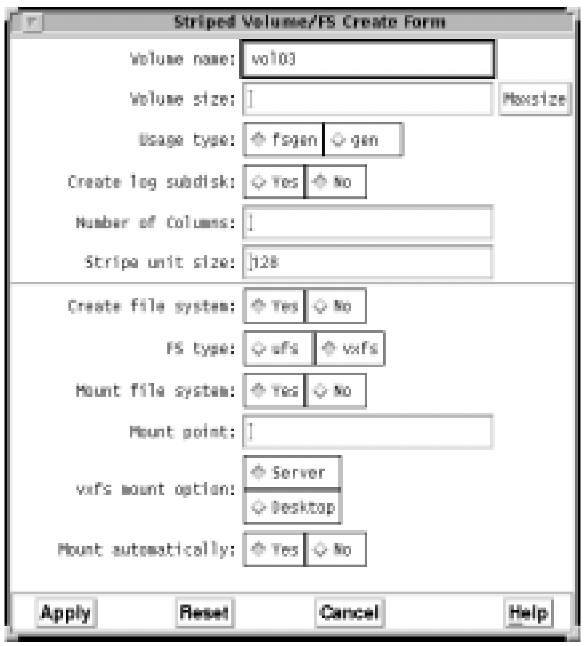

FIGURE 5-3 Striped Volume/FS Create Form

5. **Complete the Striped Volume/FS Create form.**

   For detailed information on this form, refer to online help or TABLE 5-3.

6. **When the form is properly completed, select Apply to create the volume and file system.**

   A new striped volume icon is displayed. It contains a single plex and multiple subdisks. Note that there are gaps between the subdisks of the plex to indicate that it is striped. If the file system is mounted, it is represented by the mount point, which is displayed below the new volume.

   FIGURE 5-4 shows a striped volume with a mounted file system.

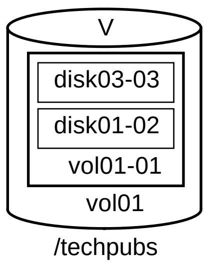

FIGURE 5-4 Striped Volume with File System

5.2.3 Creating a File System on a RAID-5 Volume

To create a file system on a RAID-5 volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Optionally, designate at least two disks (ideally, four or more) to be used for the new volume.**

   These disks must belong to the same disk group. Ensure that the disks have sufficient free space to accommodate their portion of the desired length of the new volume by clicking the RIGHT button on an unobscured portion of each disk icon to access its properties form, then checking the value in the Maximum free space field.

   If there is sufficient free space, select the disks by clicking the MIDDLE button on their icons. If no disks are selected, disks with sufficient free space are automatically used.
3. **From the Basic-Ops menu, select File System Operations, then Create.**

   A submenu listing basic volume types is displayed.
4. **Select RAID-5.**

   The RAID-5 Volume/FS Create form is displayed (see FIGURE 5-5). This form creates a file system on a RAID-5 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 file system on a new volume.

Note – Fields specific to `vxfs` file systems are only displayed if `vxfs` is installed on the system.

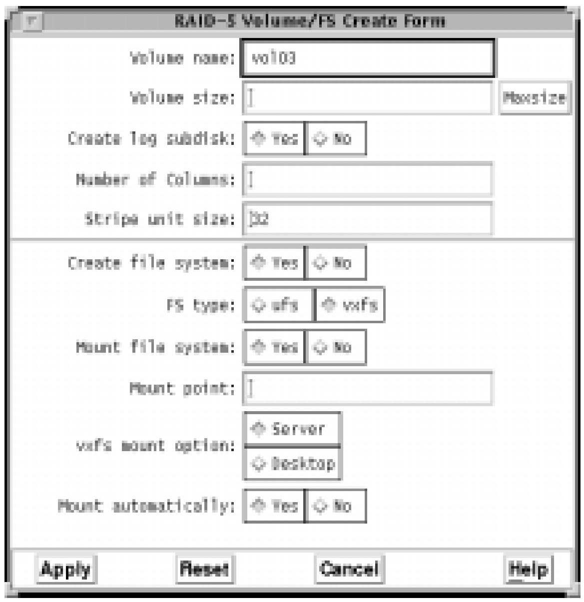

FIGURE 5-5 RAID-5 Volume/FS Create Form

5. **Complete the RAID-5 Volume/FS Create form.**

   For detailed information on this form, refer to online help or TABLE 5-4.

6. **Click Apply to create the RAID-5 volume and file system.**

   A new RAID-5 volume icon is displayed. It contains one RAID-5 plex and one log plex. If the file system is mounted, it is represented by the mount point, which is displayed below the new volume.

   FIGURE 5-6 shows a RAID-5 volume with a mounted file system, which is `/techpubs`. It also shows two plexes: one is a RAID-5 plex (`vol01-01`), and one is a log plex (`vol01-02`).

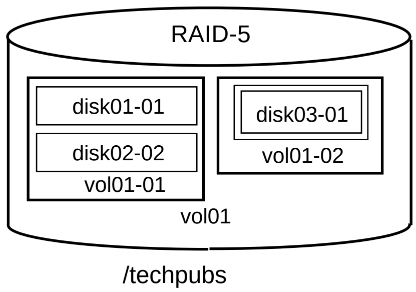

FIGURE 5-6 RAID-5 Volume with File System and Log Plex

5.2.4 Creating a File System on a Mirrored Volume

To create a file system on a mirrored volume, a volume can be created with a file system and then mirrored. The mirror layout can be simple or striped. The number of available disks must be sufficient to accommodate the layout type of both the existing volume and the mirror to be added. If a striped layout is desired, at least two disks other than those already in use by the volume must be available because one disk is required per stripe.

Note – A RAID-5 volume cannot be mirrored.

To create a file system on a mirrored volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Create a simple or striped volume containing a file system, as described in the previous sections.**
3. **Select the volume to which a mirror is to be added by clicking the LEFT button on its icon.**
4. **Optionally designate at least one disk (not already being used by the volume itself) to be used for the new mirror.**

   If the mirror is to be striped, you must select at least two disks *other* than those already in use by the volume; if the mirror is to be simple, select only one disk. These disks must belong to the same disk group.

   Ensure that the disks have sufficient free space to accommodate their portion of the desired length of the new volume by clicking the RIGHT button on an unobscured portion of each disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. For detailed information on this form, refer to TABLE 4-12. If there is sufficient free space, select the disks by clicking the MIDDLE button on their icons. If no disks are selected, disks with sufficient free space are automatically used.

5. **From the Basic-Ops menu, select Volume Operations, then Add Mirror.**

   A submenu listing mirror (plex) layout options is displayed.

6. **Select either Simple or Striped, depending on the desired mirror layout.**

   The volume's icon expands visibly and a new mirror is displayed within its borders. The new mirror layout depends on what was specified during the Add Mirror operation.

FIGURE 5-7 shows a simple, mirrored volume with a file system.

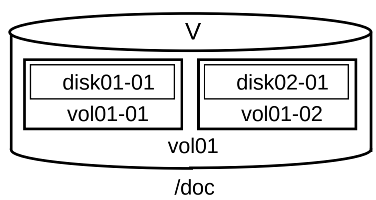

FIGURE 5-7 Mirrored Volume with File System

5.3 Creating a File System on an Existing Volume

File systems can be created and placed on existing volumes, one file system per volume. A volume icon must be selected in order for this operation to succeed.

To create a file system on an existing volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the volume on which to make the file system.**
3. **From the Basic-Ops menu, select File System Operations, then Make File System.**

   The Make File System form is displayed. Most of the form fields are already set to the defaults for making a file system on an existing volume.

4. **Complete the Make File System form (fields specific to `vxfs` file systems are only displayed if `vxfs` is installed on the system).**

   For detailed information on this form, refer to TABLE 5-5.

5. **When the form is properly completed, select Apply to create the file system.**

   If the file system is mounted, it is represented by the mount point, which is displayed below the volume icon.

5.4 Mounting a File System

A file system must be mounted before files can be accessed or created. In situations where a file system exists on a volume but is not currently mounted, that file system can be mounted at any time. A volume icon containing a valid, unmounted file system must be selected in order for this operation to succeed.

Note – The Visual Administrator has no way of knowing whether a valid, unmounted file system already exists on a given volume. You are therefore responsible for being aware of the existence of an unmounted file system on a volume, as well as that file system's characteristics.

To mount a file system on a volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the volume containing the file system to be mounted.**

3. **From the Basic-Ops menu, select File System Operations, then Mount.**

   The Mount File System form is displayed. Most of the form fields are already set to the defaults for mounting a file system on a volume.

4. **Complete the Mount File System form (fields specific to `vxfs` file systems are only displayed if `vxfs` is installed on the system).**

   For detailed information on this form, refer to TABLE 5-6.

5. **When the form is properly completed, select Apply to mount the file system.**

   When the file system is successfully mounted, it is represented by the mount point, which is displayed below the volume.

5.5 Unmounting a File System

File system administration often requires the unmounting of file systems. A file system can be unmounted from a volume as long as the mount point is not currently in use. A volume icon containing a mounted file system must be selected in order for this operation to succeed.

To unmount a file system:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the volume whose file system is to be unmounted.**
3. **From the Basic-Ops menu, select File System Operations, then Unmount.**

   The mount point should be dismissed from beneath the volume icon, indicating that the unmount operation has succeeded. If the mount point is in use, the unmount will fail.

5.6 Checking a File System for Consistency

A file system can be checked for consistency using `fsck` with the Visual Administrator. The file system must be unmounted at the time that the file system is checked. A volume icon containing an unmounted file system must be selected in order for this operation to succeed.

To check a file system for consistency:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the volume whose unmounted file system is to be checked.**
3. **From the Basic-Ops menu, select File System Operations, then Check File System (`fsck`).**
4. **The File System Check Form pops up. Click the LEFT button on Apply to start checking.**

5. **Access the Command Info Window to see if `fsck` had any problems.**

   Highlight the `fsck` command line listed in the Command History section of this window and then examine any corresponding output in the Output of the Highlighted Command section. This should reveal any `fsck`-related errors or inconsistencies. If the `fsck` line is accompanied by the word `DONE`, there are not likely to be any errors or inconsistencies present. For more information on the Command Info Window, refer to Section 3.9, “Displaying Past and Present Commands.”

Note – The Visual Administrator assumes that `fsck` will not require user input. If user input is required, the Check File System operation will fail, and `fsck` must be run from the command line.

5.7 Resizing a File System

Resizing enables a file system to be grown or shrunk according to current needs:

- A `vxfs` or `ufs` file system can be increased in size if it is no longer large enough and more space is required.
- A `vxfs` file system can be decreased in size if more space has been allocated to that file system than is actually required.

The resize operation involves resizing both the file system and its underlying 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. A volume icon containing a mounted file system must be selected in order for this operation to succeed.

Resizing can take one of four forms:

- Increase the volume and file system *to* the given length.
- Increase the volume and file system *by* the given length.
- Reduce the volume and file system *to* the given length.
- Reduce the volume and file system *by* the given length.

To resize a file system and its underlying volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the volume whose file system is to be resized.**

3. **From the Basic-Ops menu, select File System Operations, then Resize.**

   The File System Resize form is displayed (see FIGURE 5-8). Most of the form fields are already set to the defaults for resizing a file system on a volume.

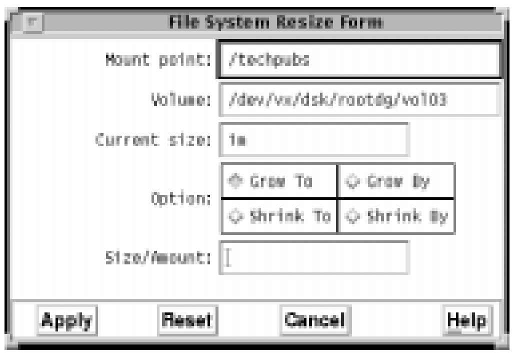

FIGURE 5-8 File System Resize Form

4. **Complete the File System Resize Form.**

   For detailed information on this form, refer to online help or TABLE 5-7.

5. **When the form is properly completed, select Apply to resize the file system.**

   As the file system and underlying volume are resized, the corresponding volume icon may change visibly.

5.8 Defragmenting a File System

Defragmentation is a technique of reorganizing a `vxfs` file system (using the `fsadm` utility). This results in improved system performance. Over time, fragmentation tends to increase on a file system, and defragmentation takes care of this.

Defragmentation can be performed on a `vxfs` file system residing on a selected volume. A volume icon containing a mounted file system must be selected in order for this operation to succeed.

Note – Defragmentation is available with `vxfs` file systems only.

Perform defragmentation on a `vxfs` file system as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the volume whose file system is to undergo defragmentation.**
3. **From the Basic-Ops menu, select File System Operations, then Defragment.**

   As defragmentation takes place, the outline of the selected volume icon momentarily disappears and then reappears.

5.9 Preparing a File System for Online Backup

The Visual Administrator can prepare a `vxfs` file system on a volume for online backup.

With the Visual Administrator, a *snapshot* image of a mounted file system is created by “snapshot” mounting another file system, which then becomes an exact copy of the first file system. The original file system is said to be *snapped*, and the copy is called the *snapshot*. The snapshot is a consistent view of the snapped file system for the point in time that the snapshot was made. For further information on the `vxfs` snapshot feature, refer to the *VERITAS File System (VxFS) System Administrator’s Guide*.

Note – The snapshot operation is available with `vxfs` file systems only.

The Visual Administrator snapshot feature can be used to create or remove a snapshot file system:

- Create — Create a simple volume and then mount a snapshot file system of the selected file system on this new volume.
- Remove — Unmount the snapshot file system and then remove the underlying volume. Removal can only be performed on a file system that was created by the Create process listed above.

A volume icon containing a mounted file system must be selected in order for either the create or remove aspect of this operation to succeed.

5.9.1 Creating a Snapshot

Create a snapshot of a file system with VxFS as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the volume whose file system is to be used for backup/snapshot purposes.**
3. **From the Basic-Ops menu, select File System Operations, then Snapshot, and then Create.**

   The FS Snapshot Create form is displayed.
4. **Complete the FS Snapshot Create form.**

   For detailed information on this form, refer to TABLE 5-8.
5. **When the form is properly completed, select Apply to take a snapshot of the file system.**

   A new, simple volume containing a mounted snapshot file system is displayed. FIGURE 5-9 shows a file system on a volume and its snapshot file system on another volume.

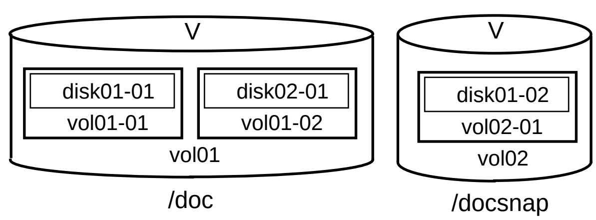

FIGURE 5-9 File System and Snapshot

The snapshot file system should now be backed up to tape (or some other media). Be careful to specify the snapshot mount point for backup.

5.9.2 Removing a Snapshot

Note – This section applies only to a snapshot file system created by the method discussed in the previous section.

Once it has served its purpose, a snapshot file system can be removed. This involves unmounting the snapshot and removing its underlying volume.

Remove a snapshot file system using the VxFS approach as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the volume containing the snapshot file system that is to be removed.**

   This volume and file system must have been created by the snapshot create operation presented in the previous section.

3. **From the Basic-Ops menu, select File System Operations, then Snapshot, and then Remove.**

   Both the snapshot file system and its underlying volume are dismissed. The resulting free space is returned to the Volume Manager free space pool.

5.10 Displaying File System Properties

The File System Properties form displays useful details on the attributes of a particular file system. A single File System Properties form provides access to all file systems known to Visual Administrator at any given time. For detailed information on this form, refer to TABLE 5-9.

Since file systems do not have associated icons, you cannot access file system properties by clicking the RIGHT button on an icon. Instead, you can access the File System Properties form in either of the following ways:

- Click the RIGHT button on the file system mount point, which is displayed below the volume icon on which the file system is mounted.
- Select Display Properties from the File System Operations menu.

FIGURE 5-10 illustrates the File System Properties form. All fields in this properties form are read-only, so this form cannot be used to alter the attributes of file systems.

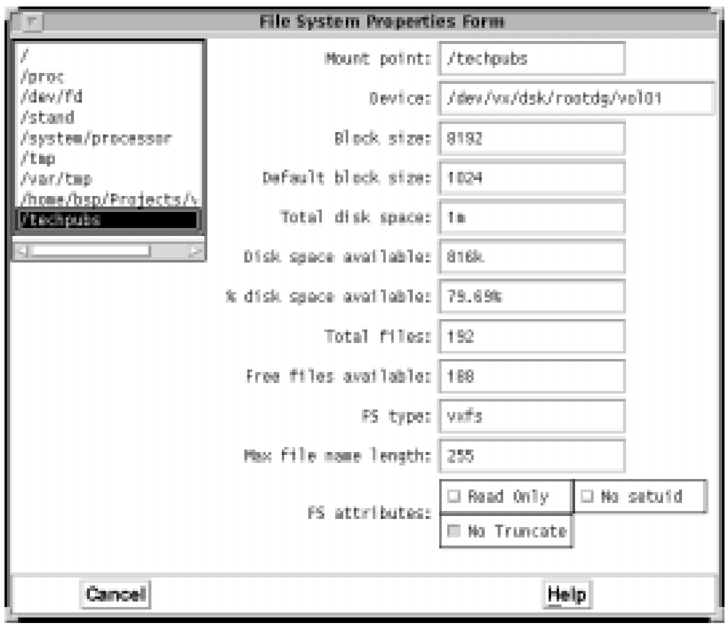

FIGURE 5-10 File System Properties Form

Display information on a specific file system with the File System Properties form as follows:

1. **Optionally select a volume whose mounted file system's properties are to be displayed.**

2. **From the Basic-Ops menu, select File System Operations, then Display Properties.**

   The File System Properties form is displayed. For detailed information on this form refer to online help or TABLE 5-9.

   If a volume has been selected, the properties for the file system that resides on that volume are displayed by default. The properties displayed on the right side of the form correspond to the selected mount point name from the list on the upper left of the form. Selecting a different mount point from the list causes the properties of that file system to be displayed instead.

5.11 Quick Reference to File System Operations Menus and Forms

This section provides information on menus and forms relating to file system operations.

Note – Some menus, forms, and form fields described here are available with the VxFS file system only. For detailed information on the VxFS file system, refer to the *VERITAS File System (VxFS) System Administrator's Guide*.

5.11.1 File System Menus

File System operations are accessed through the Basic-Ops menu. Most menus provide a Help selection, which provides information on the items and operations contained in that menu.

TABLE 5-1

Menu /Submenu Access	Description
Basic-Ops ➤ File System Operations	Provides access to File system Operations.
Basic-Ops ➤ File System Operations ➤ Create	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 may 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 Volume Manager automatically determines which disks are to be used based on available free space. From the Create menu, select the type of volume to be created from a cascading menu listing three of the basic types of volumes: Simple, Striped, or RAID-5. Requirements: * Only disks in the same disk group can be selected. * Only VM disks (disks under Volume Manager control) can be selected.
Basic-Ops ➤ File System Operations ➤ Create ➤ Simple	Creates a file system on a simple, concatenated volume whose subdisks are arranged both sequentially and contiguously within a plex (mirror). See Also: * Simple Volume/FS Create Form (TABLE 5-2)
Basic-Ops ➤ File System Operations ➤ Create ➤ Striped	Creates a file system on 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 (mirror). Requirements: * 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). RAID-5 volumes cannot be mirrored. See Also: * Striped Volume/FS Create Form (TABLE 5-3)
Basic-Ops ➤ File System Operations ➤ Create ➤ RAID-5	Creates a file system on a volume that uses striping to spread data fairly evenly across multiple disks in an array, and enables independent access. It also stripes parity across all the disks in the array. Each stripe contains a parity stripe and data stripes. Requirements: * At least three disks must be selected. See Also: * RAID-5 Volume/FS Create Form (TABLE 5-4)

TABLE 5-1 *(Continued)*

Menu /Submenu Access	Description
Basic-Ops ➤ File System Operations ➤ Create ➤ Help	Accesses a Help window that displays information relevant to the file system operations
Basic-Ops ➤ File System Operations ➤ Make File System	Makes a file system on an existing volume. You 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. See Also: * Make File System Form (TABLE 5-5)
Basic-Ops ➤ File System Operations ➤ Mount	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 be aware 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. See Also: * Mount File System Form (TABLE 5-6)
Basic-Ops ➤ File System Operations ➤ Unmount	Unmounts the file system(s) that resides on the selected volume(s). 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.
Basic-Ops ➤ File System Operations ➤ Check File System (`fsck`)	Checks the file system(s) on the selected volume(s) 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 volume(s) must contain an unmounted file system.

TABLE 5-1 *(Continued)*

Menu /Submenu Access	Description
Basic-Ops ➤ File System Operations ➤ Resize	Resizes the file system that resides on the selected volume. This involves first resizing the volume and then resizing its file system accordingly. Both the volume and file system can be increased to, increased by, reduced to, or reduced by a given length. Disk space is added to or removed from the mirrors 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. A disk cannot be selected for this operation. Requirements: * A volume icon containing a mounted file system must be selected. * The resize operation is only supported for VxFS (grow or shrink) or UFS (grow only) file systems. See Also: * File System Resize Form (TABLE 5-7)
Basic-Ops ➤ File System Operations ➤ Defragment	Performs defragmentation on the VxFS file system(s) residing on a selected volume(s). Defragmentation is a technique of reorganizing a VxFS file system (using `fsadm`). This results in improved system performance. Requirements: * At least one volume icon containing a mounted file system must be selected. * The defragmentation operation is only supported for VxFS file systems.

TABLE 5-1 *(Continued)*

Menu /Submenu Access	Description
Basic-Ops ➤ File System Operations ➤ Snapshot	Backs up a file system by creating a snapshot image of that file system. This is a convenient way of performing online backup with minimal interruption. This operation invokes the VxFS snapshot approach, which is generally preferable to the Volume Manager approach. With VxFS, a *snapshot* image of a mounted file system is created by “snapshot” mounting another file system which then becomes an exact copy of the first file system. The original file system is said to be *snapped*, and the copy is called the *snapshot*. The snapshot is a consistent view of the snapped file system for the point in time that the snapshot was made. From the Snapshot menu, a cascading menu enables you to either create or remove a snapshot. Requirements: * A volume icon containing a mounted file system must be selected in order for either the create or remove aspect of this operation to succeed (although multiple volume icons may be selected for removal). * The snapshot operation is only supported for VxFS file systems.
Basic-Ops ➤ File System Operations ➤ Snapshot ➤ Create	First creates a simple volume and then mounts a snapshot file system of the selected file system on the new volume.
Basic-Ops ➤ File System Operations ➤ Snapshot ➤ Remove	First unmounts the snapshot file system(s) and then removes the underlying volume(s). Removal can only be performed on a snapshot file system that was created by the Snapshot/Create procedure.
Basic-Ops ➤ File System Operations ➤ Display Properties	Displays information for file systems mounted on the system. You may 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.
Basic-Ops ➤ File System Operations ➤ Help	Accesses a Help window that displays information relevant to the file system operations.

5.11.2 File System Forms

Some file system operations result in the display 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.

Some form fields are required, while others are not. Fields in these forms are read/write, unless listed as read only.

Note – If vxfs is not installed when vxva is started, form fields corresponding to options that are specific to vxfs file systems are either excluded from forms or inactive.

5.11.2.1 Simple Volume/FS Create Form

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.

TABLE 5-2

Field	Description
Volume name:	Either use the default or type a different name for 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:	Type the desired volume size. The size should be typed as a number followed immediately by the letter k, m, s, or g to indicate kilobytes, megabytes, sectors, or gigabytes, 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. If you click on the Maxsize button, the Volume Manager calculates and displays the maximum volume size possible with this form's current settings. If no disks are selected, the calculation includes available space on all disks in the disk group. Otherwise, the calculation includes the selected disk(s) only.
Usage type:	Select a usage type. fsgen is the file system generic usage type that assumes the volume is being used by a file system. gen is the generic usage type that makes no assumptions regarding the data content of the volume. The default is fsgen.
Create log subdisk:	Indicate whether a log subdisk is to be created. The default is No.
Create file system:	Indicate whether a file system is to be created. When this form is invoked from the File System Operations menu, the default is to create a file system (Yes). All fields below this field are only accessible when Yes is specified here.
FS type:	Type the desired file system type.
Mount file system	Specify whether or not you want the file system to be mounted after creation. If the answer is Yes (the default), you must specify a mount point in the mount point field. All fields below this field are only accessible when Yes is specified here.

TABLE 5-2 *(Continued)*

Field	Description
Mount point:	Type the desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator will automatically create it. This field is required if the file system is to be mounted.
`vxfs` mount option:	If `vxfs` is selected as the file system type, various parameters that affect file system performance can be specified when mounting a file system. This field is used to select one of the file system configurations provided, each of which represents a certain type of environment. Mount options suitable for the chosen environment will be run. Refer to Section 5.1.2, "VxFS Mount Options," for further details on these VxFS mount options. This field is ignored if `vxfs` is not selected as the file system type.
Mount automatically:	Indicate whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/vfstab`). The default is `Yes`.

5.11.2.2 Striped Volume/FS Create Form

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.

TABLE 5-3

Field	Description
Volume name:	Either use the default or type a different name for 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:	Type the desired volume size. The size should be typed as a number followed immediately by the letter `k`, `m`, `s`, or `g` to indicate kilobytes, megabytes, sectors, or gigabytes, respectively. If no unit is specified, the default is sectors. If the size is not wholly divisible by the stripe unit size, the Volume Manager will adjust the volume size up to the next even multiple to create the volume. For a striped volume, the volume size should be calculated using: *vol_size* = *stripe_width* * *number_of_columns* * *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. If you click on the Maxsize button, the Volume Manager calculates and displays the maximum volume size possible with this form's current settings. If no disks are selected, the calculation includes available space on all disks in the disk group. Otherwise, the calculation includes the selected disk(s) only.
Usage type:	Select the desired usage type. `fsgen` is the file system generic usage type that assumes the volume is being used by a file system. `gen` is the generic usage type that makes no assumptions regarding the data content of the volume. The default is `fsgen`.

TABLE 5-3 *(Continued)*

Field	Description
Create log subdisk:	Indicate whether a log subdisk is to be created. The default is `No`.
Number of Columns:	Type the number of columns that the volume's plex (mirror) 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 columns is displayed in this field. This number corresponds to the number of disks across which data will be striped. If no number is specified, Volume Manager selects an appropriate number (usually 2).
Stripe unit size:	Type the width of the stripe units on the mirror this volume will have. The value specified may be optimized for the particular drive configuration. The default value for this field is 32 sectors, chosen as a good stripe unit size for most systems.
Create file system:	Indicate whether a file system is to be created. When this form is invoked from the File System Operations menu, the default is to create a file system (`Yes`). All fields below this field are only accessible when `Yes` is specified here.
FS type:	Type the desired file system type.
Mount file system:	Specify whether or not you want the file system to be mounted after creation. If the answer is `Yes` (the default), you must specify a mount point in the mount point field. All fields below this field are only accessible when `Yes` is specified here.
Mount point:	Type the desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator will automatically create it. This field is required if the file system is to be mounted.
`vxfs` mount option:	If `vxfs` is selected as the file system type, various parameters that affect file system performance can be specified when mounting a file system. Select one of the file system configurations provided, each of which represents a certain type of environment. Mount options suitable for the chosen environment will be run. Refer to Section 5.1.2, "VxFS Mount Options," for further details on these VxFS mount options. This field is ignored if `vxfs` is not selected as the file system type.
Mount automatically:	Indicate whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/vfstab`). The default is `Yes`.

5.11.2.3 RAID-5 Volume/FS Create Form

This form creates a file system on the new RAID-5 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.

TABLE 5-4

Field	Description
Volume name:	Either use the default or type a different name for 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:	Type the desired volume size. The size should be typed as a number followed immediately by the letter `k`, `m`, `s`, or `g` to indicate kilobytes, megabytes, sectors, or gigabytes, respectively. If no unit is specified, the default is sectors. If the size specified is not wholly divisible by the stripe unit size, it is rounded up before the volume is created. Since RAID-5 reserves separate space for parity information, the actual space contained in plexes is larger than the addressable size of the volume. The size specified in this field represents the usable space. If you click on the Maxsize button, the Volume Manager calculates and displays the maximum volume size possible with this form's current settings. If no disks are selected, the calculation includes available space on all disks in the disk group. Otherwise, the calculation includes the selected disk(s) only.
Create log subdisk:	Indicate whether a log subdisk is to be created. The default is `Yes`. At least one log subdisk should be created for a RAID-5 volume.
Number of Columns:	Type the number of columns that the volume's plex is to have. Typically, the number needed for RAID-5 should be four columns or more. When three or more VM disks are selected, this field is initialized to one less than the number of VM disks selected to enable the creation of a log subdisk.
Stripe unit size:	Indicate the width of the stripe units on the volume's plex. The value specified may be optimized for your particular drive configuration. Since 32 sectors is considered a good stripe unit size for most systems, it is specified as the default.
Create file system:	Indicate whether a file system is to be created. Since the object of this operation is to create a file system, `Yes` is the selected default.
FS type:	Type the desired file system type.
Mount file system:	Specify whether or not you want the file system to be mounted after creation. If the answer is `Yes` (the default), you must specify a mount point in the mount point field. All fields below this field are only accessible when `Yes` is specified here.
Mount point:	Type the desired mount point for the new file system. If the specified mount point does not already exist, Visual Administrator will automatically create it.
`vxfs` mount option:	If `vxfs` is selected as the file system type, various parameters that affect file system performance can be specified when mounting a file system. Select one of the file system configurations provided, each of which represents a certain type of environment. Mount options suitable for the chosen environment will be run. Refer to Section 5.1.2, "VxFS Mount Options," for further details on these VxFS mount options. This field is ignored if `vxfs` is not selected as the file system type.
Mount automatically:	Indicate whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/vfstab`). `Yes` is the default.

5.11.2.4 Make File System Form

This form is used to make a file system (using `mkfs`) according to your specifications.

TABLE 5-5

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:	Type 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:	Type the desired file system type.
Mount file system:	Specify whether or not you want the file system to be mounted after creation. If the answer is `Yes` (the default), you must specify a mount point in the mount point field. All fields below this field are only accessible when `Yes` is specified here.
Mount point:	Type the desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator will automatically create it. This field is required if the file system is to be mounted.
`vxfs` mount option:	If `vxfs` is selected as the file system type, various parameters that affect file system performance can be specified when mounting a file system. Select one of the file system configurations provided, each of which represents a certain type of environment. Mount options suitable for the chosen environment will be run. Refer to Section 5.1.2, “VxFS Mount Options,” for further details on these VxFS mount options. This field is ignored if `vxfs` is not selected as the file system type.
Mount automatically:	Indicate whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/vfstab`). `Yes` is the default.

5.11.2.5 Mount File System Form

This form is used to mount a file system that already exists on a selected volume.

TABLE 5-6

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	Type the desired file system type.
Mount point:	Type the desired mount point for the file system. If the specified mount point does not already exist, the Visual Administrator will automatically create it. The Visual Administrator attempts to provide a default mount point, which it obtains by scanning `/etc/vfstab`.
`vxfs` mount option:	If `vxfs` is selected as the file system type, various parameters that affect file system performance can be specified when mounting a file system. Select one of the file system configurations provided, each of which represents a certain type of environment. Mount options suitable for the chosen environment will be run. Refer to Section 5.1.2, "VxFS Mount Options," for further details on these VxFS mount options. This field is ignored if `vxfs` is not selected as the file system type.
Mount automatically:	Indicate whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/vfstab`). `No` is the default.

5.11.2.6 File System Resize Form

This form either grows or shrinks a file system and its underlying volume using the Volume 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. Note that a UFS file system can only be grown.

TABLE 5-7

Field	Description
Mount point:	Displays the mount point of the file system to be resized. This field is read only and cannot be changed.
Volume:	Displays the block device of the volume on which the file system resides. This field is read only and cannot be changed.

TABLE 5-7 *(Continued)*

Field	Description
Current size:	Displays the current size of the file system to be resized. This field is read only and cannot be changed.
Option:	Select 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`. The alternatives are `Grow By`, `Shrink To`, or `Shrink By`.
Size/Amount:	Type either the length to which or the amount by which the file system (and underlying 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.

5.11.2.7 FS Snapshot Create Form

This form creates a snapshot of the selected volume for backup purposes.

TABLE 5-8

Field	Description
File System:	Displays the mount point of the file system to be used as the snapshot source. This is the mount point of the file system that resides on the selected volume. This field is read only and cannot be changed.
Snapshot mount point:	Indicate the mount point of the snapshot file system to be created as a backup. The mount point must be unique
Snapshot size:	Indicate the size of the snapshot file system to be created. The snapshot file system is not an exact duplicate of the original file system, but a representation of it that does not have to be as big as the original file system. the visual Administrator calculates a default size that should be adequate for the snapshot file system. The displayed default is approximately one fifth of the length of the file system selected as the snapshot source.

5.11.2.8 File System Properties Form

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.

All fields in this form are read only and cannot be changed.

TABLE 5-9

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 the file system.
Default block size:	This indicates the smallest space allocated to a file.
Total disk space:	The amount of disk space that the file system occupies.
Disk space available:	The number of megabytes of disk storage on this file system that is available for use.
Capacity:	The percentage of the total disk storage space that is in use.
Total files:	The maximum number of files allowed on this file system.
Free files available:	The number of files that can still be created on this file system.
FS type:	The file system type (`vxfs` or `ufs`).
Max file name length:	The maximum number of characters that a file name can be on this file system. This restriction is imposed by the file system.
FS attributes:	The attributes associated with this file system: `Read Only` indicates a file system that cannot be written to. `No setuid` indicates a file system that does not support `setuid/setgid` semantics. `No Truncate` indicates a file system that does not truncate file names longer than `NAME_MAX`, the maximum file name length as defined in the system header files.

CHAPTER 6

Volume Operations

The Visual Administrator performs volume operations by executing the appropriate Volume Manager commands. This chapter provides instructions on performing the following volume operations using the Visual Administrator:

The quick reference section at the end of this chapter provides information on the Visual Administrator volume operation menus and forms.

6.1 Creating a Volume

Volumes need to be created to take advantage of the Volume Manager concept of virtual disks. Once a volume exists, a file system can be placed on the volume to organize the disk space with files and directories. Alternatively, applications like databases might be used to organize data on volumes.

The creation of a volume involves the creation of plex and subdisk components. With the basic approach to volume creation, the underlying plexes and subdisks are handled automatically. You simply indicate the desired volume characteristics.

You can create various types of volumes using either the basic or the advanced approach:

- Basic — The basic approach takes information about what the user wants to accomplish and then performs the necessary underlying tasks. This approach requires only minimal input from the user, but does permit more detailed specifications. Basic operations are accessed via the Basic-Ops menu.
- Advanced — The advanced approach consists of a number of fairly complicated commands that typically require the user to specify detailed input. These commands use a “building block” approach that requires the user to have a detailed knowledge of the underlying structure and components to manually perform the sequences of commands necessary to accomplish a certain task. Advanced operations are accessed via the Advanced-Ops menu.

Volumes can be created with various layout types:

- Simple — A concatenated volume whose subdisks are arranged both sequentially and contiguously within a plex. Concatenation enables a volume to be created from multiple regions of one or more disks if there is not enough space for an entire volume on a single region of a disk.
- Striped — A volume with data spread fairly evenly across multiple disks. *Stripe units* are equally-sized fragments that are allocated alternately and evenly to the subdisks of a single plex. There should be at least two subdisks in a striped plex, each of which should exist on a different disk. Throughput increases with the number of disks across which a plex is striped. Striping helps to balance I/O load in cases where high traffic areas exist on certain subdisks.
- RAID-5 — A volume that uses striping to spread data fairly evenly across multiple disks in an array, and enables independent access. It also stripes parity across all the disks in the array, rather than being concentrated on a single parity disk. Each stripe contains parity stripe units and data stripe units. In comparison to the performance of striped volumes, throughput of RAID-5 volumes decreases since parity information needs to be updated each time data is written. However, in comparison to mirroring, the use of parity reduces the amount of space required.
- Mirrored — A volume with multiple plexes that duplicate the information contained in a volume. Although a volume can have a single plex, at least two are required for true mirroring (redundancy of data) to be in effect. Each of these plexes should contain disk space from different disks, in order for the redundancy to be useful.
- Striped and Mirrored — A volume with a striped plex and another plex that mirrors the striped one. This requires at least two disks for striping and one or more other disks for mirroring (depending on whether the plex is simple or striped). A striped and mirrored volume is advantageous because it both spreads data across multiple disks and provides redundancy of data.

6.1.1 Creating a Simple Volume

You can create a simple volume using either a *basic* approach or an *advanced* approach.

6.1.1.1 Basic Approach

To create a simple, concatenated volume using the basic approach:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Optionally designate a disk to be used for the new volume.**

   Ensure that the disk has sufficient free space to accommodate the desired length of the new volume by clicking the RIGHT button on an unobscured portion of that disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. For more information on the VM disk properties form, refer to online help or TABLE 4-11.

   If there is sufficient free space, select the disk by clicking the LEFT button on its icon. If no disks are selected, disks with sufficient free space are automatically used.

3. **From the Basic-Ops menu, select Volume Operations, then Create.**

   A submenu listing basic volume types is displayed.

4. **Select Simple.**

   The Simple Volume/FS Create form is displayed (see FIGURE 6-1). 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 dimmed because the default is not to add a file system to the volume.

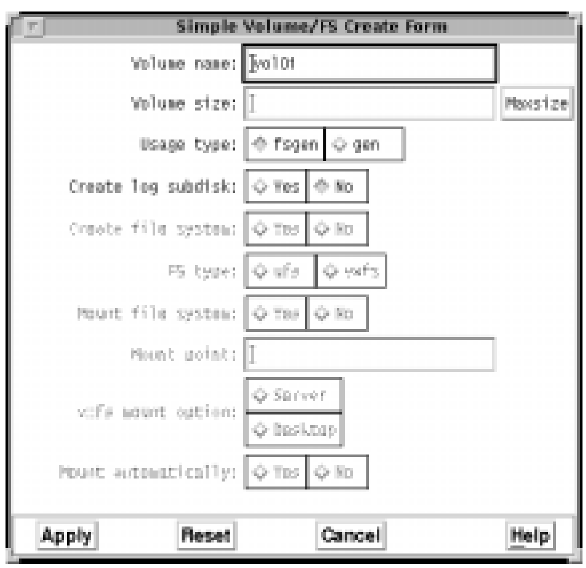

FIGURE 6-1 Simple Volume/FS Create Form

5. **Complete the Simple Volume/FS Create form.**

   For detailed information on this form, refer to online help or TABLE 6-2.

6. **When the form is properly completed, select Apply to activate the volume (and optional file system) creation.**

   A new volume icon is displayed. Since this is a simple, concatenated volume, it contains a single plex. If a file system exists and is mounted, it is represented by the mount point, which is displayed below the new volume.

   FIGURE 6-2 illustrates a simple, concatenated volume.

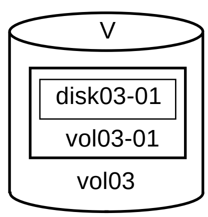

FIGURE 6-2 Simple Volume

6.1.1.2 Advanced Approach

To create a simple, concatenated volume using the advanced approach:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Designate the disk on which to create the subdisk.**

   Ensure that the disk has sufficient free space to accommodate the desired length of the new volume by clicking the RIGHT button on an unobscured portion of that disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. If there is sufficient free space, select the disk by clicking the LEFT button on its icon. For more information on the VM disk properties form, refer to online help or TABLE 4-11.

3. **From the Advanced-Ops menu, select Subdisk, then Create.**

   The Subdisk Create form is displayed. This form establishes the parameters for the subdisk. Some of the fields already contain default values and others are optional; only the subdisk length is required.

4. **Complete the Subdisk Create form, typing the subdisk length and optionally altering other fields.**

   Type the length of the subdisk to be created. 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. The maximum length of this field is 14 characters.

   For more detailed information on this form, refer to online help or TABLE 8-2.

5. **When the form is properly completed, select Apply to activate the subdisk creation.**

   A new subdisk icon is displayed in the selected disk icon.

6. **Drag the new subdisk icon (by pressing and continuing to hold the LEFT button) beyond the borders of its disk and into an open area of the view window. Drop the subdisk by releasing the LEFT button.**

   The Plex Create form is displayed. When this form is invoked in this way, it will automatically create a new plex and associate the dragged subdisk with it. The Plex Create form establishes the parameters for the plex. Some of the fields already contain default values and others are optional. Since this is to be a concatenated (simple) volume, the layout must be set to Concatenated.

7. **Complete the Plex Create form, indicating a concatenated layout and optionally altering any fields.**

   For more detailed information on this form, refer to online help or TABLE 7-2.

8. **When the form is properly completed, select Apply to activate both the plex creation and the association of subdisk and plex.**

   A plex icon containing a subdisk icon is displayed. The fact that the subdisk icon is within the plex icon confirms their association. The length of the plex should be the same as that of the subdisk.

9. **Select the new plex by clicking the LEFT button on its icon.**

   When the plex is selected immediately prior to the creation of the volume in this way, the plex will automatically be associated with the new volume.

10. **From the Advanced-Ops menu, select Volume, then Create.**

    The Volume Create form is displayed. This form establishes the parameters for the volume. Some of the fields already contain default values and others are optional. Notice that the length is the same as that of the subdisk in the selected plex.

11. **Complete the Volume Create form, optionally altering any fields.**

    For more detailed information on this form, refer to online help or TABLE 6-7.

12. **When the form is properly completed, select Apply to activate both the volume creation and the association of plex and volume.**

    A volume icon is displayed around the plex and its associated subdisk. The fact that the plex icon is within the volume icon confirms their association.

13. **Initialize the volume:**

    a. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

    b. **Select the volume that is to be initialized by clicking the LEFT button on its icon.**

    c. **From the Advanced-Ops menu, select Volume, then Initialize.**

    A submenu is displayed and offers the following choices for the volume initialization: Active, Enable, Clean, or Zero.

    d. **Select Active (this is recommended for most circumstances).**

6.1.2 Creating a Striped Volume

You can create a striped volume using either the *basic* approach or the *advanced* approach.

6.1.2.1 Basic Approach

To create a striped volume using the basic approach:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Optionally designate at least two disks to be used for the new volume.**

   These disks must belong to the same disk group. Ensure that the disks have sufficient free space to accommodate their portion of the desired length of the new volume by clicking the RIGHT button on an unobscured portion of each disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. For more detailed information on this form, refer to online help or TABLE 4-11.

   If there is sufficient free space, select the disks by clicking the MIDDLE button on their icons. If no disks are selected, disks with sufficient free space are automatically used.

3. **From the Basic-Ops menu, select Volume Operations, then Create.**

   A submenu listing basic volume types is displayed.

4. **Select Striped.**

   The Striped Volume/FS Create form is displayed (see FIGURE 6-3). This form creates a striped 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 dimmed because the default is not to add a file system to the volume.

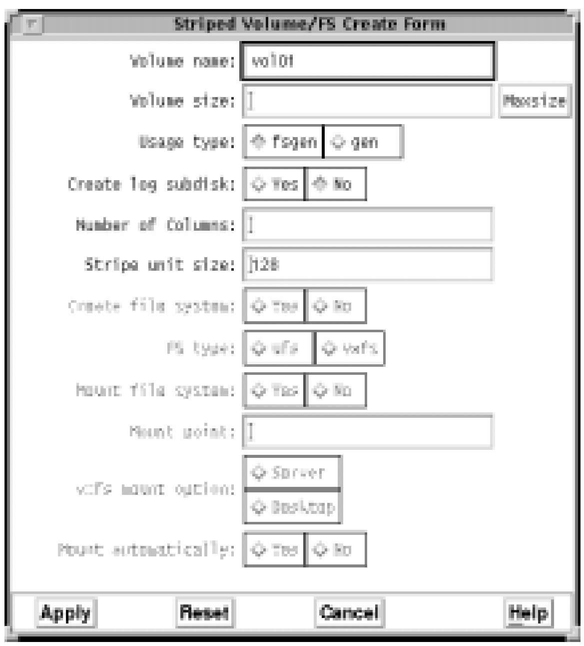

FIGURE 6-3 Striped Volume/FS Create Form

5. **Complete the Striped Volume/FS Create form.**

   For detailed information, refer to online help or TABLE 6-3.

6. **When the form is properly completed, select Apply to activate the volume and file system creation.**

   A new volume icon is displayed. Since this is a striped volume, it contains a single plex and multiple subdisks. Note that there are gaps between the subdisks of the plex to indicate that it is striped. If a file system exists and is mounted, it is represented by the mount point, which is displayed below the new volume.

   FIGURE 6-4 shows a striped volume.

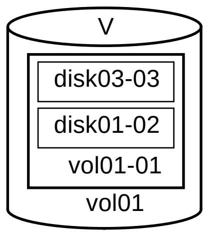

FIGURE 6-4 Striped Volume

6.1.2.2 Advanced Approach

When creating a striped volume, at least two disks must be selected across which to stripe the data on the volume's plex.

To create a striped volume with two subdisks on different disks:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Designate the disk on which to create the first subdisk.**

   Ensure that the disk has sufficient free space to accommodate the desired length of the new volume by clicking the RIGHT button on an unobscured portion of that disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. If there is sufficient free space, select the disk by clicking the LEFT button on its icon. For more detailed information on this form, refer to online help or TABLE 4-11.
3. **From the Advanced-Ops menu, select Subdisk, then Create.**

   The Subdisk Create form is displayed. This form establishes the parameters for the subdisk. Some of the fields already contain default values and others are optional; only the subdisk length is required.
4. **Complete the Subdisk Create form, typing the subdisk length and optionally altering other fields.**

   For detailed information on this form, refer to online help or TABLE 8-2.
5. **When the form is properly completed, select Apply to activate the subdisk creation.**

   A new subdisk icon is displayed in the selected disk icon.
6. **Repeat the previous steps to select another disk in this disk group and create a subdisk of the same length as the first one.**

7. **Select both of the new subdisks by clicking the MIDDLE button on their icons.**
8. **From the Advanced-Ops menu, select Plex, then Create.**

   The Plex Create form is displayed. When this form is invoked while subdisks are selected, it will automatically create a new plex and associate the selected subdisks with it. The Plex Create form establishes the parameters for the plex. Some of the fields already contain default values and others are optional. The number of stripes is set to the number of selected subdisks (two). Since this is to be a striped volume, the layout must be set to Striped and the stripe unit size must be typed.

9. **Complete the Plex Create form, indicating stripe information and optionally altering any other fields.**

   For detailed information on this form, refer to online help or TABLE 7-2.

10. **When the form is properly completed, select Apply to activate both the plex creation and the association of subdisks and plex.**

    A plex icon containing both subdisk icons is displayed. The fact that the subdisk icons are within the plex icon confirms their association.

11. **Select the plex by clicking the LEFT button on its icon.**

    When the plex is selected immediately prior to the creation of the volume in this way, it will automatically be associated with the new volume.

12. **From the Advanced-Ops menu, select Volume, then Create.**

    The Volume Create form is displayed. This form establishes the parameters for the volume. Some of the fields already contain default values and others are optional. Notice that the length is the combined length of the subdisks in the selected plex.

13. **Complete the Volume Create form, optionally altering any fields.**

    For detailed information on this form, refer to online help or TABLE 6-7.

14. **When the form is properly completed, select Apply to activate both the volume creation and the association of plexes and volume.**

    A volume icon is displayed around the plex and its associated subdisks. The fact that the plex icon is within the volume icon confirms their association.

15. **Initialize the volume:**

    a. **Select the volume that is to be initialized by clicking the LEFT button on its icon.**

    b. **From the Advanced-Ops menu, select Volume, then Initialize.**

       A submenu is displayed and offers four choices for the volume initialization: Active, Enable, Clean, or Zero.

    c. **Select Active (this is recommended for most circumstances).**

6.1.3 Creating a RAID-5 Volume

To create a RAID-5 volume using the basic approach:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Optionally, designate at least three disks (ideally, four or more) to be used for the new volume.**

   These disks must belong to the same disk group. Ensure that the disks have sufficient free space to accommodate their portion of the desired length of the new volume by clicking the RIGHT button on an unobscured portion of each disk icon to access its properties form, then checking the value in the Maximum free space field.

   If there is sufficient free space, select the disks by clicking the MIDDLE button on their icons. If no disks are selected, disks with sufficient free space are automatically used.

3. **From the Basic-Ops menu, select Volume Operations, then Create.**

   A submenu listing basic volume types is displayed.

4. **Select RAID-5.**

   The RAID-5 Volume/FS Create form is displayed (see FIGURE 6-5). This form creates a RAID-5 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 dimmed because the default is to add a file system to the volume only when you access this menu via the Basic Ops, File Systems Operations menu.

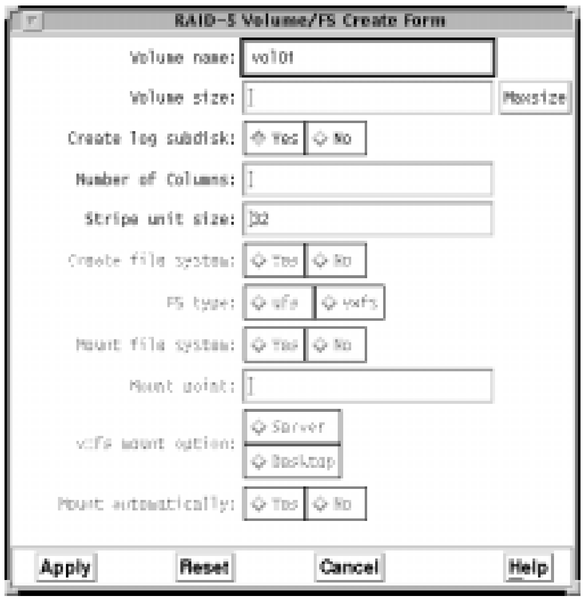

FIGURE 6-5 RAID-5 Volume / FS Create Form

5. **Complete the RAID-5 Volume/FS Create form.**

   For detailed information on this form, refer to online help or TABLE 6-4.

6. **When the form is properly completed, select Apply to activate the volume and file system creation.**

   A new volume icon is displayed. Since this is a RAID-5 volume, it contains one RAID-5 plex with multiple subdisks and one log plex. Note that there are gaps between the subdisks of the RAID-5 plex to indicate that it is striped. The double outline indicates the log plex. If a file system exists and is mounted, it is represented by the mount point, which is displayed below the new volume.

   FIGURE 6-6 shows a RAID-5 volume.

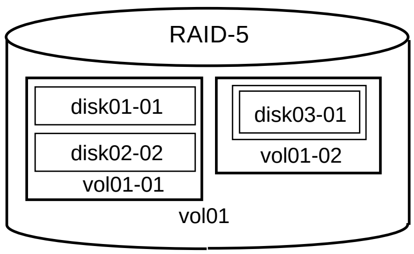

FIGURE 6-6 RAID-5 Volume

6.1.4 Creating a Mirrored Volume

A mirrored volume has multiple plexes that duplicate the information contained in a volume. Although a volume can have a single plex, at least two plexes on different disks are required for true mirroring (redundancy of data) to be in effect.

To create a mirrored volume, a volume must be created and then mirrored. The mirror layout can be simple or striped. The number of available disks must be sufficient to accommodate the layout type of both the existing volume and the mirror to be added. If a striped layout is desired, at least two disks other than those already in use by the volume must be available because one disk is required per stripe.

You can create a mirrored volume using either the *basic* approach or the *advanced* approach.

6.1.4.1 Basic Approach

To create a mirrored volume using the basic approach:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Create a simple volume, as described in a previous section.**
3. **Select the new volume to which a mirror is to be added by clicking the LEFT button on its icon.**

4. **Optionally designate at least one disk (not already being used by the volume itself) to be used for the new mirror.**

   If the mirror is to be striped, at least two disks *other* than those already in use by the volume must be selected; if the mirror is to be simple, only one disk is selected. These disks must belong to the same disk group. Ensure that the disks have sufficient free space to accommodate their portion of the desired length of the new volume by clicking the RIGHT button on an unobscured portion of each disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. For details on this form, refer to TABLE 4-11.

   If there is sufficient free space, select the disks by clicking the MIDDLE button on their icons. If no disks are selected, disks with sufficient free space are automatically used.

5. **From the Basic-Ops menu, select Volume Operations, then Add Mirror.**

   A submenu listing mirror layout options is displayed.

6. **Select either Simple or Striped, depending on the desired mirror layout.**

   The volume's icon expands visibly and a new mirror is displayed within its borders. The new mirror layout depends on what was specified during the Add Mirror operation.

6.1.4.2 Advanced Approach

When creating a mirrored volume, at least two plexes must be created on at least two disks. Both plexes must then be associated with the same volume.

To create a mirrored volume with two plexes, each containing a single subdisk:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Designate the disk on which to create the first subdisk.**

   Ensure that the disk has sufficient free space to accommodate the desired length of the new volume by clicking the RIGHT button on an unobscured portion of that disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. If there is sufficient free space, select the disk by clicking the LEFT button on its icon. For details on this form, refer to TABLE 4-11.

3. **From the Advanced-Ops menu, select Subdisk, then Create.**

   The Subdisk Create form is displayed. This form establishes the parameters for the subdisk. Some of the fields already contain default values and others are optional; only the subdisk length is required.

4. **Complete the Subdisk Create form, typing the subdisk length and optionally altering other fields.**

   For details on this form, refer to online help or TABLE 8-2.

5. **When the form is properly completed, select Apply to activate the subdisk creation.**

   A new subdisk icon is displayed in the selected disk icon.

6. **Drag the new subdisk icon (by pressing and continuing to hold the LEFT button) beyond the borders of its disk and into an open area of the view window. Drop the subdisk by releasing the LEFT button.**

   The Plex Create form is displayed. When this form is invoked in this way, it will automatically create a new plex and associate the dragged subdisk with it. This form establishes the parameters for the plex. Some of the fields already contain default values and others are optional. Since this volume is to contain two concatenated (simple) plexes for mirroring, the layout must be set to Concatenated.

7. **Complete the Plex Create form, indicating a concatenated layout and optionally altering any other fields.**

   For details on this form, refer to online help or TABLE 7-2.

8. **When the form is properly completed, select Apply to activate both the plex creation and the association of subdisk and plex.**

   A plex icon containing a subdisk icon is displayed. The fact that the subdisk icon is within the plex icon confirms their association. The length of the plex should be the same as that of the subdisk.

9. **Repeat the previous steps to select another disk in this disk group, create a subdisk of the same length as the first one on that disk, create a plex of the same type as the first, and associate the second subdisk with the second plex.**

10. **Select both plexes by clicking the MIDDLE button on their icons.**

    When the plexes are selected immediately prior to the creation of the volume in this way, they will automatically be associated with the new volume.

11. **From the Advanced-Ops menu, select Volume, then Create.**

    The Volume Create form is displayed. This form establishes the parameters for the volume. Some of the fields already contain default values and others are optional. Notice that the length is the same as that of the subdisks in the selected plexes.

12. **Complete the Volume Create form, optionally altering any fields.**

    For details on this form, refer to online help or TABLE 6-7.

13. **When the form is properly completed, select Apply to activate both the volume creation and the association of plexes and volume.**

    A volume icon is displayed around both plexes and their associated subdisks. The plexes are side-by-side in the volume, to illustrate that they are mirrored. The fact that both plex icons are within the volume icon confirms their association.

14. **Initialize the volume:**

    a. **To initialize the new volume, click the LEFT button on its icon.**

    b. **From the Advanced-Ops menu, select Volume, then Initialize. A submenu is displayed and offers four choices for the volume initialization: Active, Enable, Clean, or Zero.**

    c. **Select Active (this is recommended for most circumstances).**

6.1.5 Creating a Striped and Mirrored Volume

A striped and mirrored volume is advantageous because it both spreads data across multiple disks and provides redundancy of data.

Although there is currently no menu item for creating a volume that is both striped and mirrored, it is possible to do so. This is done by creating a striped volume and then adding a mirror (in the form of either a simple or striped plex) to that volume.

The number of available disks must be sufficient to accommodate the layout type of both the existing volume and the mirror to be added. At least three disks are required to create a striped, mirrored volume. The striped plex requires two disks (one per stripe); the plex that acts as the mirror must occupy a third disk if it is simple or a third and fourth disk if it is also striped. If any stripes or mirrors occupy the same disk, the striping and mirroring will not be effective.

You can create a striped and mirrored volume using either a *basic* approach or an *advanced* approach.

6.1.5.1 Basic Approach

To create a striped and mirrored volume using the basic approach:

1. **Follow the steps for creating a striped volume, as described previously.**

2. **Select the new volume to which a mirror is to be added by clicking the LEFT button on its icon.**

3. **Optionally designate at least one disk (not already being used by the volume itself) to be used for the new mirror.**

   If the mirror is to be striped, at least two disks *other* than those already in use by the volume must be selected; if the mirror is to be simple, only one disk is selected. These disks must belong to the same disk group. Ensure that the disks have sufficient free space to accommodate their portion of the desired length of the new volume by clicking the RIGHT button on an unobscured portion of each disk icon to access its VM disks properties form, then checking the value in the Maximum free space field. For details on this form, refer to TABLE 4-11.

   If there is sufficient free space, select the disks by clicking the MIDDLE button on their icons. If no disks are selected, disks with sufficient free space are automatically used.

4. **From the Basic-Ops menu, select Volume Operations, then Add Mirror.**

   A submenu listing mirror layout options is displayed.

5. **Select either Simple or Striped, depending on the desired mirror layout.**

   The striped volume's icon expands visibly and a new mirror is displayed within its borders. The new mirror layout depends on what was specified during the Add Mirror operation.

   FIGURE 6-7 illustrates a striped volume with a simple mirror.

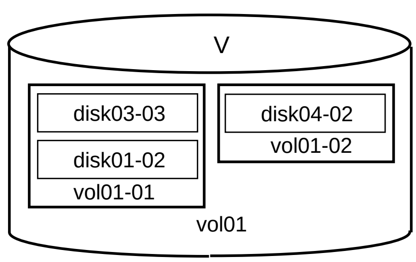

FIGURE 6-7 Striped and Mirrored Volume

6.1.5.2 Advanced Approach

The simplest method is to create a striped volume and then add a mirror to that volume. Since a striped volume is being mirrored, it is preferable for the mirror layout to be striped to maintain the advantages of striping. At least four disks are required to create a striped volume with a striped mirror; the original striped plex requires two disks and the plex that acts as a striped mirror requires an additional two disks for mirroring to be effective. If only three disks are available, a simple mirror can be added instead (although this causes some performance drawbacks).

To create a striped and mirrored volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Follow the steps for creating a striped volume, as described in the previous section.**

3. **Designate a third disk on which to create the first subdisk for the striped mirror.**

   Projection can be used to ensure that the third disk is not currently being used by the striped volume's subdisks (by pressing Shift-RIGHT on the striped volume's icon and observing which disks now contain highlighted subdisks). The disk should have sufficient space to accommodate a subdisk as long as the subdisk length in the striped volume. If there is sufficient free space, select the disk by clicking the LEFT button on its icon.

4. **From the Advanced-Ops menu, select Subdisk, then Create.**

   The Subdisk Create form is displayed. This form establishes the parameters for the subdisk. Some of the fields already contain default values and others are optional; only the subdisk length is required.

5. **Complete the Subdisk Create form, typing the subdisk length (that matches that of an existing subdisk in the striped volume) and optionally altering other fields.**

   For details on this form, refer to online help or TABLE 8-2.

6. **When the form is properly completed, select Apply to activate the subdisk creation.**

   A new subdisk icon is displayed in the selected disk icon.

7. **Create an identically-sized subdisk on a fourth disk.**

   Using the LEFT button, drag and drop the newly-created subdisk icon onto a disk icon containing sufficient space to accommodate a subdisk of the same length. Make sure the disk is not already in use by the striped volume or the selected subdisk.

   A dialog box is displayed to announce that the dragged subdisk was not associated and that Visual Administrator has therefore created a new subdisk rather than moving one. Select the Continue button to proceed. A new subdisk icon is displayed on its targeted disk.

8. **Select both of the new subdisks by clicking the MIDDLE button on their icons.**

9. **From the Advanced-Ops menu, select Plex, then Create.**

   The Plex Create form is displayed. This form establishes the parameters for the plex. Some of the fields already contain default values and others are optional. The number of stripes is set to the number of selected subdisks (two). For a striped plex, the layout must be set to Striped.

10. **Complete the Plex Create form, indicating stripe information and optionally altering any other fields.**

    For details on this form, refer to online help or TABLE 7-2.

11. **When the form is properly completed, select Apply to activate both the plex creation and the association of subdisk and plex.**

    A plex icon containing two subdisk icons is displayed. The fact that the subdisk icons are within the plex icon confirms their association. The length of the plex should be the same as that of the volume to be mirrored.

12. **The new plex icon can now be associated with the striped volume, thereby putting mirroring into effect. Drag and drop the new plex icon onto the striped volume icon. Position the plex so that the pointer is directly over an unobscured portion of the volume icon.**

    The striped volume's icon expands visibly and the mirror is displayed within its borders.

13. **Initialize the volume:**

    a. **Select the volume that is to be initialized by clicking the LEFT button on its icon.**

    b. **From the Advanced-Ops menu, select Volume, then Initialize. A submenu is displayed and offers four choices for the volume initialization: Active, Enable, Clean, or Zero.**

    c. **Select Active (this is recommended for most circumstances).**

6.2 Removing a Volume

Volumes may be removed when they are no longer needed or when the disk space that they occupy needs to be reused.

The removal of a volume involves the removal of plex and subdisk components. With the basic approach to volume removal, the underlying plexes and subdisks are handled automatically. You simply indicate which volume is to be removed; the volume and its components are then removed and the disk space set aside for that volume is de-allocated. A volume must be selected for this operation to succeed.

Caution – 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.

You can remove a volume using either a *basic* approach or an *advanced* approach.

6.2.1 Basic Approach

To remove a volume using the basic approach:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the volume to be removed by clicking the LEFT button on its icon.**
3. **From the Basic-Ops menu, select Volume Operations, then Remove Volumes Recursively.**

   If the selected volume is enabled and therefore in danger of losing valuable data, a Visual Administrator Warning window is displayed to announce this fact. The warning window requires input before Visual Administrator will proceed with anything. If certain that the data on the selected volume does not need to be preserved, select OK to proceed with the removal; otherwise, select Cancel to abandon the removal.

   The volume is dismissed when the removal is complete.

6.2.2 Advanced Approach

When a volume is removed via the Advanced-Ops menu, its plex and subdisk components are not automatically removed; it is necessary to manually remove each of its components. It is generally not a good idea to leave the subdisks of a removed volume behind, as the space occupied by these subdisks cannot be used by the Volume Manager's free space pool until they are removed.

Note – A volume must be stopped before it can be removed. Refer to the section on stopping volumes for information on how to stop a volume.

If a volume is to be removed instantly and recursively, the easiest approach is to use the Basic-Ops menu's volume removal option. However, the following procedure describes how to manually remove a volume and some or all of its components.

To remove a volume with a single plex and subdisk:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Stop the volume, as described in the section entitled "Stopping Volumes."**

3. **Select the volume to be removed by clicking the LEFT button on its icon.**
4. **From the Advanced-Ops menu, select Volume, then Remove Volumes.**

   The volume icon momentarily shrinks and is displayed next to its plex icon, then disappears entirely. The volume is now removed, but its plex and subdisk remain.

   To remove multiple volumes simultaneously, follow the previous steps, but use the MIDDLE button to select multiple volumes for removal at once.

   To remove a removed volume's plex:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of** `rootdg`**, by default).**
2. **Select the plex to be removed by clicking the LEFT button on its icon.**
3. **From the Advanced-Ops menu, select Plex, then Remove Plexes.**

   The plex icon momentarily shrinks and is displayed next to its previous position, then disappears entirely along with the subdisk icon. The plex is now removed, but its subdisk remains on its disk and is now categorized as a free subdisk.

   To remove multiple plexes simultaneously, follow the previous steps, but use the MIDDLE button to select multiple plexes for removal at once.

   To remove a removed volume and plex's subdisk:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of** `rootdg`**, by default).**
2. **Identify the subdisk icon that formerly belonged to the removed plex and volume.**

   If projection is set to show free subdisks (via the Projection menu), then this subdisk icon should be highlighted on its disk.

3. **Select the subdisk to be removed by clicking the LEFT button on its icon.**
4. **From the Advanced-Ops menu, select Subdisk, then Remove Subdisks.**

   The subdisk icon is dismissed, leaving a gap (or "hole") on its disk.

   To remove multiple subdisks simultaneously, follow the previous steps, but use the MIDDLE button to select multiple subdisks for removal at once.

6.3 Adding a Mirror to a Volume

A volume is mirrored to provide redundancy of data. Each plex on a mirrored volume duplicates the information stored on that volume.

A mirror can be added to an existing volume by associating another plex (of the correct length) to the volume. The layout of the mirror to be added can be either simple or striped. The number of available disks must be sufficient to accommodate the layout type of both the existing volume and the mirror to be added.

A volume must be selected for this operation to succeed. Disks that are not already in use by the volume itself can be selected for this operation.

To add a mirror to an existing volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the volume to which a mirror is to be added by clicking the LEFT button on its icon.**

3. **Optionally designate at least one disk (not already being used by the volume itself) to be used for the new mirror.**

   If the mirror is to be striped, at least two disks *other* than those already in use by the volume must be selected; if the mirror is to be simple, only one disk is selected. These disks must belong to the same disk group. Ensure that the disks have sufficient free space to accommodate their portion of the desired length of the new volume by clicking the RIGHT button on an unobscured portion of each disk icon to access its VM disk properties form, then checking the value in the Maximum free space field. For detailed information on this form, refer to TABLE 4-11.

   If there is sufficient free space, select the disks by clicking the MIDDLE button on their icons. If no disks are selected, disks with sufficient free space are automatically used.

4. **From the Basic-Ops menu, select Volume Operations, then Add Mirror.**

   A submenu listing mirror layout options is displayed.

5. **Select either Simple or Striped, depending on the desired mirror layout.**

   The volume's icon expands visibly and a new mirror is displayed within its borders. The new mirror layout depends on what was specified during the Add Mirror operation.

FIGURE 6-8 illustrates a simple, mirrored volume.

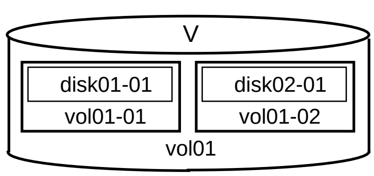

FIGURE 6-8 Mirrored Volume

6.3.1 Adding a Mirror to a Root Volume

When a volume is selected for mirroring, Visual Administrator automatically checks whether that volume is a root volume (containing bootable information). If it is, a dialog box is displayed and asks whether to make the new mirror bootable.

To add a mirror to a root volume:

1. **Go to the View of `rootdg` (which is the only view that can contain a root volume).**

2. **Select the root volume to which a mirror is to be added by clicking the LEFT button on its icon.**

   In the Visual Administrator, a root volume is displayed as a normal volume icon (named `rootvol`, by default) with a mounted file system named `/`.

3. **Follow the steps for adding a mirror to a volume, as described in the previous section.**

4. **When the Visual Administrator detects that the volume to be mirrored is a root volume, a dialog box is displayed and asks whether to make the new mirror bootable. The response to this should be `Yes` so that the mirror's disk can be used for booting if the root volume's original disk fails.**

   If the root volume's mirror is made bootable, mirrors are also added to the volume `swapvol` (but not to any other volumes that happen to have subdisks that reside on the root disk) and slicing is maintained.

6.4 Adding a Log

This operation adds a log plex containing a log subdisk to a volume. For mirrored striped or concatenated volumes, this will add a Dirty Region Log. For RAID-5, it will add a RAID-5 log.

Note – Logs are only useful for mirrored volumes or RAID-5 volumes.

To add a log:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select a simple, striped, or RAID-5 volume by clicking the LEFT button on the icon.**

3. **From the Basic-Ops menu, select Volume Operations, then Add Log.**

A log associated with a subdisk will be displayed in the volume icon. A simple or striped DRL log subdisk or RAID-5 log can be identified by their double borders. Logs are useful in the event of disk failures.

6.5 Removing a Mirror or Log

A mirror or log can be removed recursively, which involves automatically removing any subdisks associated with that plex. Both, mirrors and logs, that are associated with volumes, and those that are not, can be removed in this way. Space formerly allocated to the removed subdisks is returned to the free space pool.

Note – The last valid mirror in a started or enabled volume cannot be removed, unless the volume is stopped before the remove operation.

To remove a mirror or log recursively:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the mirror or log to be removed by clicking the LEFT button on its icon.**
3. **From the Basic-Ops menu, select Volume Operations, then Remove Mirror or Log.**

   The mirror or log and any associated subdisk icons are dismissed, indicating that these objects have been removed. Though a RAID-5 log can be removed by the same operation, it is advisable to keep at least one log associated with the RAID-5 volume to help cope with disk failures.

6.6 Resizing a Volume

An existing volume can be resized when this becomes necessary. Resizing enables a volume to be grown or shrunk according to current needs. A volume icon must be selected in order for this operation to succeed.

Note – If a `vxfs` or `ufs` file system resides on a volume, the resize operation offered through the File System Operations menu should be used to adjust both the file system and the underlying volume.

The volume size may need to be increased if a file system or application on the volume requires additional space. 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. When shrinking a volume, it is important to be aware that the volume may contain data that could be lost.

The Visual Administrator does not permit a volume containing a mounted file system to be shrunk via the volume resize operation offered through the Basic-Ops menu. This prevents the possibility of losing the contents of the end of a file system whose volume has been shrunk.

Note – Since the Visual Administrator cannot tell whether an unmounted file system resides on a volume, shrinking a volume with an unmounted file system may result in unrecoverable data loss.

Resizing can take one of four forms:

- Increase the volume *to* the given length.
- Increase the volume *by* the given length.
- Reduce the volume *to* the given length.
- Reduce the volume *by* the given length.

To resize a volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the volume whose file system is to be resized.**

3. **From the Basic-Ops menu, select Volume Operations, then Resize.**

   The Volume Resize form is displayed (see FIGURE 6-9). Most of the form fields are already set to the defaults for resizing a volume.

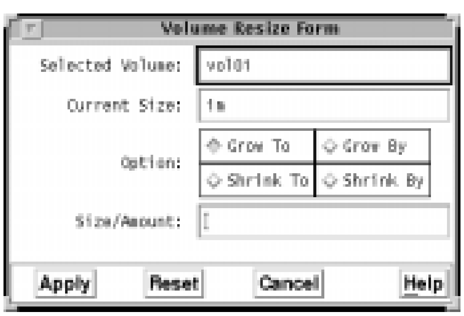

FIGURE 6-9 Volume Resize Form

4. **Complete the Volume Resize form.**

   For detailed information on this form, refer to online help or TABLE 6-5.

5. **When the form is properly completed, select Apply to activate the resize operation.**

   As the volume is resized, the corresponding volume icon may change visibly.

6.7 Preparing a Volume for Online Backup

The Visual Administrator provides access to two forms of online backup, each of which functions differently:

- The VxFS snapshot approach, offered through the File System Operations menu
- The Volume Manager snapshot approach, offered through the Volume Operations menu

This section focuses on the Volume Manager approach to preparing a volume for online backup. Details on the VxFS approach (which is often preferable) are available in the chapter on file system operations.

With the Volume Manager, the snapshot operation creates a new volume that is a snapshot of an existing volume. This is done by creating a mirror of the existing volume (creating and associating a plex) using disk space from the pool of free disk space. The mirror is brought up to date (this may take some time) and a separate (snapshot) volume is then created for it. 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 at a convenient time without stopping it. After the backup is made, the snapshot volume can be removed without losing any data.

The Volume Manager uses the `vxassist` utility operations to handle online backup of volumes with minimal interruption.

The `vxassist snapstart` operation creates a write-only backup mirror, which is attached to and synchronized with the volume to be backed up. The amount of synchronization time depends on the size of the volume. When synchronized with the volume, the backup mirror is ready to be used as a snapshot mirror. However, it continues being updated until it is detached during the actual snapshot portion of the procedure.

The online backup procedure is completed by running the `vxassist snapshot` command on the volume with the snapshot mirror. This operation detaches the finished snapshot mirror, creates a new normal volume, and attaches the snapshot mirror to it. The snapshot then becomes a read-only volume.

The amount of time involved in creating the snapshot mirror is long in contrast to the brief amount of time that it takes to create the snapshot volume.

6.7.1 Creating a Snapshot

To create a snapshot of a volume using the Volume Manager approach:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the volume to be used for backup/snapshot purposes.**
3. **From the Basic-Ops menu, select Volume Operations, then Snapshot, then Snapstart.**

   A snapshot mirror is displayed within the selected volume. As the snapshot mirror is updated, its icon is dimmed. This synchronization takes a variable amount of time, which could be significant (depending on the size of the selected volume).
4. **Once the snapshot mirror is fully updated (and its icon is no longer dimmed), prepare to make the snapshot volume by optionally asking users to save files and temporarily reduce activity.**

   It is also a good idea to sync file systems at this point.
5. **Select the volume containing the snapshot mirror.**
6. **From the Basic-Ops menu, select Volume Operations, then Snapshot, then Snapshot again.**

   The Snapshot form is displayed.
7. **Complete the Snapshot form.**

   For detailed information on this form, refer to online help or TABLE 6-6.
8. **When the form is properly completed, select Apply to activate the snapshot operation.**

   When the backup plex is synchronized with the mirror, the `vxassist snapshot` operation begins. This portion of the procedure should only take a brief amount of time. The snapshot operation detaches the finished snapshot plex, creates a new normal volume, and attaches the snapshot plex to it. As this occurs, the original volume icon returns to its former state and the backup plex moves over into the new snapshot volume (which has the name specified in the Snapshot name field of the Snapshot form). The new snapshot volume remains as a consistent copy of the selected volume at the time the snapshot was begun.
9. **Back up the snapshot volume to tape (or some other media).**

10. **Remove the snapshot volume when it is no longer necessary, as it takes up as much space as a normal volume.**

    FIGURE 6-10 illustrates a volume and its snapshot.

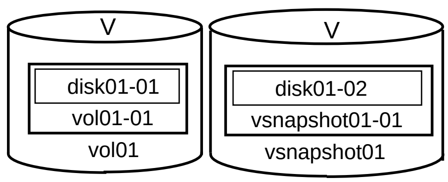

FIGURE 6-10 Volume and Snapshot

6.7.2 Removing a Snapshot

Once it has served its purpose, a snapshot volume can be removed in the same way as a regular volume is removed.

To remove a snapshot volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the snapshot volume to be removed.**
3. **From the Basic-Ops menu, select Volume Operations, then Remove Volumes Recursively.**

   The snapshot volume and its components are dismissed after prompting for confirmation.

6.8 Starting and Stopping Volumes

At times, it may be necessary to start or stop a volume. Starting involves enabling a disabled volume, while stopping involves disabling an enabled one. The Visual Administrator sometimes notifies you that a volume needs to be started before a desired operation can be performed.

Volumes can switch between the following states: Started, Unstartable, and Startable.

Volumes should be initialized before they are started and must have been started before they can be stopped.

6.8.1 Starting Volumes

A volume must be started before most operations can be performed on that volume. With the Visual Administrator, either a single volume or all volumes can be started at one time.

To start a single volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Before starting the volume, confirm that it is not started.**

   To do this, access its properties form by clicking RIGHT on its icon. The Utility State field should indicate that the volume is unstartable (never been started) or startable (currently, not started).
3. **Select the volume to be started by clicking the LEFT button on its icon.**
4. **From the Advanced-Ops menu, select Volume, then Start Volumes, and then Start.**

   The volume icon should rewrite itself as it is being started.
5. **Confirm that the volume has now been started by accessing its properties form again.**

   The Utility State field should now indicate that the volume is started.

To start all volumes simultaneously:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **From the Advanced-Ops menu, select Volume, then Start Volumes, and then Start All.**

   All volume icons that are not started should rewrite themselves as they are being started. Volumes that were already started are not affected.
3. **Check the properties forms of any of the newly started volumes to confirm that they are indeed started.**

6.8.2 Stopping Volumes

A volume must be stopped before some operations (such as removal of the volume itself or dissociation of its last plex) can be performed on that volume. A volume containing a mounted file system cannot be stopped. With Visual Administrator, either a single volume or all volumes can be stopped at one time.

To stop a single volume:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Before stopping the volume, confirm that it is started.**

   To do this, access its properties form by clicking RIGHT on its icon. The Utility State field should indicate that the volume is started.
3. **Select the volume to be stopped by clicking the LEFT button on its icon.**
4. **From the Advanced-Ops menu, select Volume, then Stop Volumes, and then Stop.**

   The volume icon should rewrite itself as it is being stopped.
5. **Confirm that the volume has now been stopped by accessing its properties form again.**

   The Utility State field should now indicate that the volume is startable.

To stop all volumes simultaneously:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **From the Advanced-Ops menu, select Volume, then Stop Volumes, and then Stop All.**

   All started volume icons should rewrite themselves as they are being stopped. Volumes that were already stopped are not affected. For details on the Volume Operations menu choices, refer to TABLE 6-1.
3. **Check the properties forms of any of the newly stopped volumes to confirm that they are indeed startable.**

6.9 Displaying and Altering Properties

The properties of the volume corresponding to a volume icon can be displayed (and potentially adjusted) in a properties form. The Volume Properties form (illustrated in FIGURE 6-11) provides detailed information about the characteristics of a particular volume.

To access the properties form for a particular volume icon:

- **Move the pointer onto that icon and click the RIGHT button. (If the volume is under analysis, press Shift-RIGHT instead.)**

  The Volume Properties form is displayed and displays details of the volume's configuration. For detailed information on this form, refer to TABLE 6-8.

Some of the displayed properties (such as the volume's name) can be altered directly through this form by editing the appropriate field(s) and then selecting the form's Apply button. Selecting the Reset button rereads the volume properties and displays the current values.

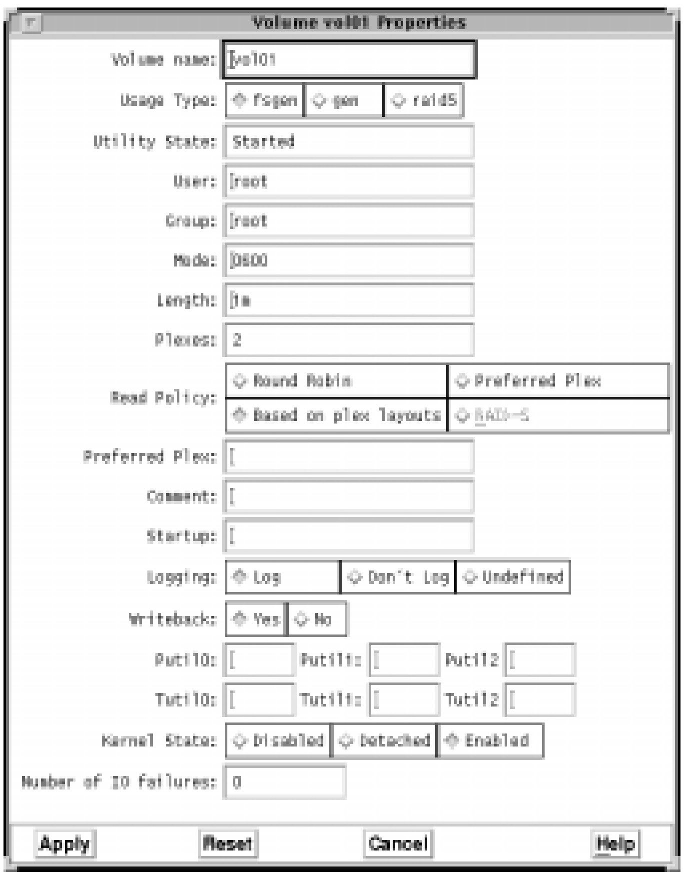

FIGURE 6-11 Volume Properties Form

6.10 Quick Reference to Volume Operations Menus and Forms

This section provides information on volume menus and forms.

6.10.1 Volume Menus

Both the Basic-Ops and Advanced-Ops menus provide access to volume-related menus. Most menus provide a Help selection, which provides information on items and operations contained in that menu.

TABLE 6-1

Menu/Submenu Access	Description
Basic-Ops ➤ Volume Operations ➤ Create	Creates a volume from one or more disks. You may select one or more disks on which to create the volume (providing there is sufficient space on the disks). If no disks are specified, the Volume Manager automatically determines which disks are to be used based on available free space. From the Create menu, select the type of volume to be created from a cascading menu listing three of the basic types of volumes: Simple, Striped, or RAID-5. If a mirrored volume is desired, a simple or striped volume must be created and then mirrored using the Add Mirror option. Requirements: * Only disks in the same disk group can be selected. * Only VM disks (disks under Volume Manager 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. See Also: * Simple Volume/FS Create Form (TABLE 6-2) * Striped Volume/FS Create Form (TABLE 6-3) * RAID-5 Volume/FS Create Form (TABLE 6-4)
Basic-Ops ➤ Volume Operations ➤ Create ➤ Simple	Creates a simple, concatenated volume whose subdisks are arranged both sequentially and contiguously within a plex (mirror).
Basic-Ops ➤ Volume Operations ➤ Create ➤ 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.
Basic-Ops ➤ Volume Operations ➤ Create ➤ RAID-5	Creates a volume that uses striping to spread data fairly evenly across multiple disks in an array, and enables independent access. It also stripes parity across all the disks in the array. Each stripe contains a parity stripe and data stripes.

TABLE 6-1 *(Continued)*

Menu/Submenu Access	Description
Basic-Ops ➤ Volume Operations ➤ Remove Volumes Recursively	Removes the selected volume(s) and deallocates all of the disk space set aside for that volume. It automatically removes all underlying mirrors (plexes) and subdisks associated with the volume. *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 (i.e., started or enabled).* Requirements: * At least one volume icon must be selected. * The selected volume(s) cannot contain a mounted file system.
Basic-Ops ➤ Volume Operations ➤ Add Mirror	Adds a mirror to the selected volume by associating a mirror (plex) of the correct length to the volume. The mirror effectively duplicates the information contained in the volume. Although a volume can have a single mirror, at least two are required for true mirroring (redundancy of data) to be in effect. From the Add Mirror menu, select the type of mirror to be added from a cascading menu listing two of the basic types of mirrors: Simple and Striped. Disks can be selected for this operation. However, the number of selected disks must be sufficient to accommodate the layout type of both the existing volume and the mirror to be added. If no disks are selected, the free space for the mirror is allocated by the Volume Manager. *Note: You cannot mirror a RAID-5 volume.* Requirements: * A volume icon must be selected. * For a striped mirror, at least two disks other than those already in use by the volume must be available.
Basic-Ops ➤ Volume Operations ➤ Add Mirror ➤ Simple	Adds a simple, concatenated mirror whose subdisks are arranged both sequentially and contiguously.
Basic-Ops ➤ Volume Operations ➤ Add Mirror ➤ Striped	Adds a mirror whose data is allocated evenly across each of its subdisks in an alternating fashion. This is accomplished with *stripes*, which are relatively small, equally-sized fragments that are allocated alternately to each subdisk.
Basic-Ops ➤ Volume Operations ➤ Add Log	Adds a log plex containing a log subdisk to a volume. For striped and concatenated volumes, this adds a Dirty Region Log (DRL). For RAID-5, this operation adds a RAID-5 log.

TABLE 6-1 *(Continued)*

Menu/Submenu Access	Description
Basic-Ops ➤ Volume Operations ➤ Remove Mirror or Log	Removes the selected mirror (plex) or log, along with any associated subdisks. Requirements: * A mirror or log icon must be selected. * The last valid mirror in a started or enabled volume cannot be removed, unless the volume is stopped before the remove operation.
Basic-Ops ➤ Volume Operations ➤ Resize	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/from the plexes associated with the volume. If new disk space is needed during the resize, it is allocated as necessary. Any unused space is added to the free space pool. Requirements: * A volume icon must be selected. * A volume containing a mounted file system cannot be shrunk. See Also: * Volume Resize Form (TABLE 6-5)
Basic-Ops ➤ Volume Operations ➤ Snapshot	Backs up a volume by creating a snapshot image of that volume. This is a convenient way of performing online backup with minimal interruption. This operation invokes the Volume Manager snapshot approach, in which the snapshot operation creates a new volume that is a snapshot of an existing volume. This is done by creating a mirror of the existing volume (creating and associating a plex) using disk space from the pool of free disk space. The mirror is brought up to date (this may take some time) and a separate (snapshot) volume is then created for that mirror. 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, the snapshot volume can be removed without losing any data. From the Snapshot menu, a cascading menu enables you to first create the snapshot mirror and then the snapshot volume: Snapstart and Snapshot. Requirements: * A volume icon must be selected. * There must be sufficient free disk space to accommodate the snapshot volume. See Also: * Snapshot Form (TABLE 6-6)

TABLE 6-1 *(Continued)*

Menu/Submenu Access	Description
Basic-Ops ➤ Volume Operations ➤ Snapshot ➤ Snapstart	Starts the snapshot procedure by creating a snapshot mirror within the volume to be backed up. It takes a variable amount of time to update the new mirror, during which time the snapshot mirror icon is ghosted.
Basic-Ops ➤ Volume Operations ➤ Snapshot ➤ Snapshot	At a convenient time (preferably after warning users to reduce activity briefly), create another volume for the snapshot mirror. This portion of the procedure should take only seconds to complete.
Basic-Ops ➤ Volume Operations ➤ Help	Accesses a Help window that displays information relevant to the basic volume operations.
Advanced-Ops ➤ Volume ➤ Create	Creates a volume. You may select one or more plexes (mirrors) to be associated with the new volume after creation. See Also: * Volume Create Form (TABLE 6-7)
Advanced-Ops ➤ Volume ➤ Remove Volumes	Removes the selected volume(s). If the selected volume is started, it must be stopped before it can be removed. *This is a permanent operation and cannot be undone. Any plexes (mirrors) 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.
Advanced-Ops ➤ Volume ➤ Initialize Volumes	Initializes the selected volume(s). From the Initialize Volumes menu, select the type of initialization from a cascading menu listing four choices: Active, Enable, Clean, and Zero. 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 (mirror) that is complete (or contiguous).
Advanced-Ops ➤ Volume ➤ Remove Volumes ➤ Active	Enables the selected volume and its associated plexes (mirrors), and sets the state of all associated plexes to ACTIVE.
Advanced-Ops ➤ Volume ➤ Remove Volumes ➤ Enable	Enables the selected volume and its associated plexes (mirrors), but leaves the plex states as EMPTY.
Advanced-Ops ➤ Volume ➤ Remove Volumes ➤ Clean	Sets the state for all associated plexes (mirrors) of the selected volume to CLEAN. This can be applied only under limited circumstances.

TABLE 6-1 *(Continued)*

Menu/Submenu Access	Description
Advanced-Ops ➤ Volume ➤ Remove Volumes ➤ Zero	Enables the selected volume and its associated plexes (mirrors), 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.
Advanced-Ops ➤ Volume ➤ Start Volumes	Starts the selected volume(s). A volume must be started before it can be accessed. From the Start volumes menu, a cascading menu enables you to indicate whether all volumes or just those selected should 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 must be initialized before it can be started.
Advanced-Ops ➤ Volume ➤ Start Volumes ➤ Start	Starts the selected volume, which must be startable.
Advanced-Ops ➤ Volume ➤ Start Volumes ➤ Start All	Starts all volumes in this disk group that can be started.
Advanced-Ops ➤ Volume ➤ Stop Volumes	Stops the selected volume(s). A stopped volume is inaccessible. From the Stop volumes menu, a cascading menu enables you to indicate whether all volumes or just those selected should be stopped 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.
Advanced-Ops ➤ Volume ➤ Stop Volumes ➤ Stop	Stops the selected volume.
Advanced-Ops ➤ Volume ➤ Stop Volumes ➤ Stop All	Stops all volumes in this disk group.

TABLE 6-1 *(Continued)*

Menu/Submenu Access	Description
Advanced-Ops ➤ Volume ➤ Resynchronize Volumes	Brings all plexes (mirrors) within the selected volume(s) up to date. Any plexes that are inconsistent are resynchronized to contain consistent data. Depending on how current the plexes are, this operation may take some time. Requirements: * At least one volume icon must be selected. * The selected volume(s) must be started.
Advanced-Ops ➤ Volume ➤ Set to Maint State	Sets the state of the selected volume(s) to a maintenance state. Refer to the `vxvol`(1M) manual page for information on the maintenance state. Requirements: * At least one volume icon must be selected.
Advanced-Ops ➤ Volume ➤ Help	Accesses a Help window that displays information relevant to the basic volume operations.

6.10.2 Volume Forms

Some volume operations result in the display 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.

Some form fields are required, while others are not. Fields in these forms are read / write unless listed as read only.

Note – If `vxfs` is not installed when `vxva` is started, form fields corresponding to options that are specific to `vxfs` file systems are either excluded from forms or dimmed.

6.10.2.1 Simple Volume/FS Create Form

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 dimmed because the default is not to add a file system to the volume.

TABLE 6-2

Field	Description
Volume name:	Either use the default or type a different name for 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:	Type the desired volume size. The size should be typed as a number followed immediately by the letter `k`, `m`, `s`, or `g` to indicate kilobytes, megabytes, sectors, or gigabytes, 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. If you click on the Maxsize button, the Volume Manager calculates and displays the maximum volume size possible with this form’s current settings. If no disks are selected, the calculation includes available space on all disks in the disk group. Otherwise, the calculation includes the selected disk(s) only.
Usage Type:	Select the desired usage type. `fsgen` is the file system generic usage type that assumes the volume is being used by a file system. `gen` is the generic usage type that makes no assumptions regarding the data content of the volume. The default is `fsgen`.
Create log subdisk:	Indicate whether a log subdisk is to be created. The default is `No`.
Create file system:	Indicate whether a file system is to be created. When this form is invoked from the Volume Operations menu, the default is not to create a file system (`No`). All fields below this field are only accessible when `Yes` is specified here.
FS type:	Select the desired file system type.
Mount file system:	Specify whether or not you want the file system to be mounted after creation. If the answer is `Yes` (the default), you must specify a mount point in the mount point field. All fields below this field are only accessible when `Yes` is specified here.
Mount point:	Type the desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator will automatically create it. This field is required if the file system is to be mounted.
`vxfs` mount option:	If `vxfs` is selected as the file system type, various parameters that affect file system performance can be specified when mounting a file system. Select one of the file system configurations provided, each of which represents a certain type of environment. Mount options suitable for the chosen environment will be run. This field is ignored if `vxfs` is not selected as the file system type.
Mount automatically:	Indicate whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/vfstab`). The default is `Yes`.

6.10.2.2 Striped Volume/FS Create Form

This form creates a striped 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 dimmed because the default is not to add a file system to the volume.

TABLE 6-3

Field	Description
Volume name:	Either use the default or type a different name for 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:	Type the desired volume size. The size should be typed as a number followed immediately by the letter `k`, `m`, `s`, or `g` to indicate kilobytes, megabytes, sectors, or gigabytes, respectively. If no unit is specified, the default is sectors. If the size is not wholly divisible by the stripe unit size, Volume Manager will adjust the volume size up to the next even multiple to create the volume. For a striped volume, the volume size should be calculated as follows: *vol_size = stripe_width * number_of_columns * 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. If you click on the Maxsize button, the Volume Manager calculates and displays the maximum volume size possible with this form's current settings. If no disks are selected, the calculation includes available space on all disks in the disk group. Otherwise, the calculation includes the selected disk(s) only.
Usage Type:	Select the desired usage type. `fsgen` is the file system generic usage type that assumes the volume is being used by a file system. `gen` is the generic usage type that makes no assumptions regarding the data content of the volume. The default is `fsgen`.
Create log subdisk:	Indicate whether a log subdisk is to be created. The default is `No`.
Number of Columns:	Type the number of columns that the volume's plex (mirror) 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 columns is displayed in this field. This number corresponds to the number of disks across which data will be striped. If no number is specified, Volume Manager selects an appropriate number (usually 2).
Stripe unit size:	Indicate the width of the stripe units on the plex (mirror) that this volume will have. The value specified may be optimized for the particular drive configuration. The default value for this field is 32 sectors, chosen as a good stripe unit size for most systems.
Create file system:	Indicate whether a file system is to be created. When this form is invoked from the Volume Operations menu, the default is not to create a file system (`No`). All fields below this field are only accessible when `Yes` is specified here.
FS type:	Select the desired file system type.

TABLE 6-3 *(Continued)*

Field	Description
Mount file system:	Specify whether or not you want the file system to be mounted after creation. If the answer is `Yes` (the default), you must specify a mount point in the mount point field. All fields below this field are only accessible when `Yes` is specified here.
Mount point:	Type the desired mount point for the new file system. If the specified mount point does not already exist, the Visual Administrator will automatically create it. This field is required if the file system is to be mounted.
`vxfs` mount option:	If `vxfs` is selected as the file system type, various parameters that affect file system performance can be specified when mounting a file system. Select one of the file system configurations provided, each of which represents a certain type of environment. Mount options suitable for the chosen environment will be run. This field is ignored if `vxfs` is not selected as the file system type.
Mount automatically:	Indicate whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/vfstab`). The default is `Yes`.

6.10.2.3 RAID-5 Volume / FS Create Form

This form creates a RAID-5 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 dimmed because the default is not to add a file system to the volume.

TABLE 6-4

Field	Description
Volume name:	Either use the default or type a different name for 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:	Type the desired volume size. The size should be typed as a number followed immediately by the letter `k`, `m`, `s`, or `g` to indicate kilobytes, megabytes, sectors, or gigabytes, respectively. If no unit is specified, the default is sectors. If the size specified is not wholly divisible by the stripe unit size, it is rounded up before the volume is created. Since RAID-5 reserves separate space for parity information, the actual space contained in plexes is larger that the addressable size of the volume. The size specified in this field represents the usable space. If you click on the Maxsize button, the Volume Manager calculates and displays the maximum volume size possible with this form’s current settings. If no disks are selected, the calculation includes available space on all disks in the disk group. Otherwise, the calculation includes the selected disk(s) only.
Create log subdisk:	Indicate whether a log subdisk is to be created. The default is `Yes`. At least one log subdisk should be created for a RAID-5 volume.

TABLE 6-4 *(Continued)*

Field	Description
Number of columns:	Type the number of columns that the volume's plex is to have. Typically, the number needed for RAID-5 should be four columns or more.
Stripe unit size:	Indicate the width of the stripe units on the volume's plex. The value specified may be optimized for your particular drive configuration. Since 32 sectors is considered a good stripe unit size for most systems, it is specified as the default.
Create file system:	Indicate whether a file system is to be created. Since the object of this operation is to create a file system, `Yes` is the selected default.
FS type:	Select the desired file system type.
Mount file system:	Specify whether or not you want the file system to be mounted after creation. If the answer is `Yes` (the default), you must specify a mount point in the mount point field. All fields below this field are only accessible when `Yes` is specified here.
Mount point:	Type the desired mount point for the new file system. If the specified mount point does not already exist, Visual Administrator will automatically create it.
`vxfs` mount option:	If `vxfs` is selected as the file system type, various parameters that affect file system performance can be specified when mounting a file system. Select one of the file system configurations provided, each of which represents a certain type of environment. Mount options suitable for the chosen environment will be run. This field is ignored if `vxfs` is not selected as the file system type.
Mount automatically:	Indicate whether this file system should be mounted every time the system comes up (by placing an entry in `/etc/vfstab`). `Yes` is the default.

6.10.2.4 Volume Resize Form

This form either grows or shrinks a volume using the Volume 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.

TABLE 6-5

Field	Description
Selected Volume:	Displays the name of the volume to be resized. This field is read only and cannot be changed.

TABLE 6-5 *(Continued)*

Field	Description
Current size:	Displays the current size of the volume to be resized. This field is read only and cannot be changed.
Option:	Select 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:	Type 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.

6.10.2.5 Snapshot Form

This form creates a snapshot of the selected volume for backup purposes.

TABLE 6-6

Field	Description
Selected Volume:	Displays the name of the volume to be used as the snapshot source. This field is read only and cannot be changed. There must be sufficient free space to accommodate the snapshot volume.
Snapshot name:	Type the name of the snapshot volume to be created as a backup. Although a default name is displayed in this field, a name that more closely resembles that of the selected volume should be used for easier association. The maximum length is 14 characters. The snapshot name must be unique.

6.10.2.6 Volume Create Form

This form creates a volume according to your specifications.

TABLE 6-7

Field	Description
Volume name:	Either use the default or type a different name for 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:	Select the desired usage type. `fsgen` is the file system generic usage type that assumes the volume is being used by a file system. `gen` is the generic usage type that makes no assumptions regarding the data content of the volume. The default is `fsgen`. This field is optional.
User:	Type 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 14 characters.
Group:	Type 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 14 characters.
Mode:	Set 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 4 characters.
Length:	Type 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:	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:	Displays the read policy 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. RAID-5 – Indicates a RAID-5 layout.
Preferred Plex:	Lists the name of the preferred plex (mirror) 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:	enables you to add an appropriate comment for this volume. The maximum length of the comment is 40 characters. This field is optional.
Startup:	May 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.

TABLE 6-7 *(Continued)*

Field	Description
Logging:	Indicates whether logging is defined and supported on this volume. An undefined log type is included to support old versions of the Volume 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 (mirror). The default is `No`.
Putil0:	Lists the permanent utility field 0. This is reserved for Volume Manager use, but may be changed. The maximum length of all Putil fields is 14 characters. This field is optional.
Putil1:	Lists the permanent utility field 1. This field is reserved, but may be changed. This field is optional.
Putil2:	Lists the permanent utility field 2. This field is reserved, but may be changed. This field is optional.

6.10.2.7 Volume Properties Form

This form provides detailed information on the attributes of a particular volume. 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 the Apply button.

Click RIGHT mouse button on desired volume icon. (If volume icon is undergoing analysis, press Shift- RIGHT instead.)

TABLE 6-8

Field	Description
Volume name:	The default or type a different name for 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 typing another name in this field.
Usage Type:	Volume usage type. `fsgen` is the file system generic usage type that assumes the volume is being used by a file system. `gen` is the generic usage type that makes no assumptions regarding the data content of the volume. `root` and `swap` volumes cannot be created and are presented for informational purposes only. These usage types have restrictions on possible operations. A `root` volume contains a root file system, while a `swap` volume is used as the system swap device. This field is read only and cannot be changed.
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 14 characters.

TABLE 6-8 *(Continued)*

Field	Description
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 14 characters.
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 4 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:	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 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. RAID-5 – Indicates a RAID-5 layout.
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:	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 Volume 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 (mirror).
Putil0:	The permanent utility field 0. This is reserved for Volume Manager use, but may be changed. The maximum length of all Putil fields is 14 characters.
Putil1:	The permanent utility field 1. This field is reserved, but may be changed.
Putil2:	The permanent utility field 2. This field is reserved, but may be changed.
Tutil0:	The temporary utility fields 0. This is reserved for Volume Manager use, but can be changed. The maximum length of all Tutil fields is 14 characters.
Tutil1:	The temporary utility field 1. This field is reserved, but can be changed.

TABLE 6-8 *(Continued)*

Field	Description
Tutil2:	The temporary utility field 2. This field is reserved, but can be changed.
Kernel State:	The accessibility of the volume. This field is read only and cannot be changed.
Number of IO failures:	The number of I/O operations that have failed on this volume. This field is read only and cannot be changed.

Plex Operations

The Visual Administrator performs plex operations by executing the appropriate Volume Manager commands. This chapter provides instructions on performing the following plex operations using the Visual Administrator:

The quick reference section at the end of this chapter provides information on the Visual Administrator plex operations menus and forms.

FIGURE 7-1 illustrates plex associated with a volume (`vol03-01`), a dissociated plex (`pl-01`), and a dissociated plex containing an associated subdisk (`vol01-01`).

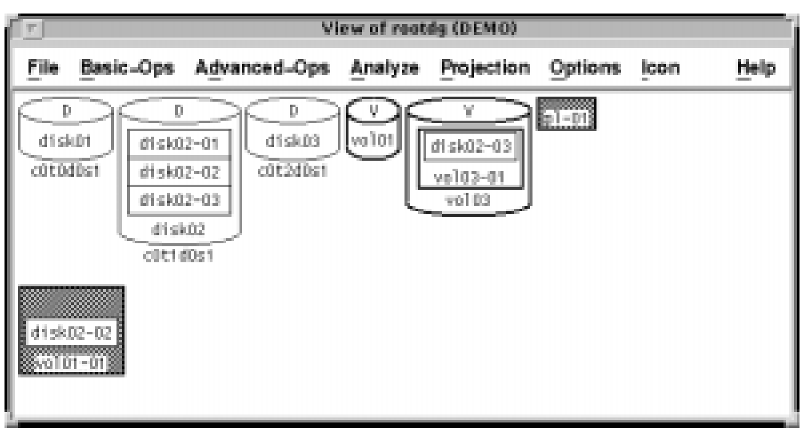

FIGURE 7-1 Plex Associations

7.1 Creating a Plex

Plexes need to be created to connect subdisks to volumes. Without plexes, volumes are useless.

The creation of a plex alone is not particularly useful unless one or more subdisks are associated with that plex. If one or more subdisks have been selected at the time of the plex creation, they will automatically be associated with the new plex. A striped plex requires at least two subdisks.

Create a plex and associate it with a single subdisk as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Create a new subdisk, as described in Chapter 8, "Subdisk Operations."**
3. **Drag the new subdisk icon beyond the borders of its disk and drop it into an open area of the view window.**

   The Plex Create form is displayed. When this form is invoked in this way, it will automatically create a new plex and associate the dragged subdisk with it. The Plex Create form establishes the parameters for the plex. Some of the fields already contain default values and others are optional.
4. **Complete the Plex Create form according to the type of plex to be created.**

   For detailed information on this form, refer to online help or TABLE 7-2.
5. **When the form is properly completed, select Apply to activate both the plex creation and the association of subdisk and plex.**

   A plex icon containing a subdisk icon is displayed. The fact that the subdisk icon is within the plex icon confirms their association.

Create a plex and associate it with multiple (two) subdisks as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Create two new subdisks, as described in Chapter 8, "Subdisk Operations."**
3. **Select both of the new subdisks by clicking the MIDDLE button on their icons.**
4. **From the Advanced-Ops menu, select Plex, then Create.**

   The Plex Create form is displayed. When this form is invoked while subdisks are selected, it will automatically create a new plex and associate the selected subdisks with it. The Plex Create form establishes the parameters for the plex. Some of the fields already contain default values and others are optional.

5. **Complete the Plex Create form according to the type of plex to be created.**

   If the plex is to be striped, type a stripe width and the number of columns. For detailed information on this form, refer to online help or TABLE 7-2.

6. **When the form is properly completed, select Apply to activate both the plex creation and the association of subdisks and plex.**

   A plex icon containing both subdisk icons is displayed. The fact that the subdisk icons are within the plex icon confirms their association.

7.2 Associating a Plex

Once a plex has been created, that plex can be associated with a volume. A volume is of little use until it has at least one associated plex. When multiple plexes are associated with a volume, mirroring is in effect. If the volume is already started and an additional plex is being associated, the Volume Manager will update the newer plex by copying all necessary data to that plex; this operation may take some time.

Several plexes can be associated with a single volume at a time, but only one volume can be selected per association operation.

There are a few methods of associating a plex with a volume:

- Selecting the volume and plex(es) and then using menus.
- Dragging each plex icon onto the volume icon and then dropping the plex icon on that volume when the pointer (in the image of a hand) is positioned over an unobscured portion of the volume icon.
- Automatically associating one or more plexes with a volume by selecting the plex icon(s) immediately prior to the creation of a volume. See Chapter 6, “Volume Operations,” for further details.

Associate a plex with an existing volume via menus as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select both the plex and the volume to which it is to be associated by clicking the MIDDLE button on both of their icons.**

3. **From the Advanced-Ops menu, select Plex, then Associate Plexes.**

   The volume and plex icons are dismissed momentarily and then the plex icon is displayed within the volume icon. The plex is now associated with the volume.

   To associate multiple plexes with a single volume, select multiple plexes along with the volume and repeat the previous procedure.

7.3 Dissociating a Plex

Associated plexes can be dissociated from their parent volumes. However, the last plex in a started volume cannot be dissociated until that volume is stopped. A plex may be dissociated if it is no longer required to reflect a copy of a volume. Plexes must be dissociated before they can be removed.

There are a couple of methods of dissociating a plex from a volume:

- Selecting the plex(es) and then using menus.
- Dragging each plex icon beyond the borders of the volume icon and then dropping the plex in an open area of the view window.

Dissociate a plex from its volume via menus as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the plex to be dissociated from its volume by clicking the LEFT button on its icon.**

3. **From the Advanced-Ops menu, select Plex, then Dissociate Plexes.**

   The plex icon is dismissed momentarily and then is displayed beyond the borders of the volume icon. The plex is now dissociated from its volume.

   To dissociate multiple plexes from a volume, select multiple plexes (using the MIDDLE button) and repeat the previous procedure.

7.4 Detaching a Plex

A detached plex is inaccessible for reads and writes, but is still associated with a volume. One or more plexes can be detached from their parent volume at a time. The volume must be started before a plex can be detached.

Note – This operation is not permitted when the specified plex is the last valid plex on the volume.

Detach a plex as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the plex to be detached by clicking the LEFT button on its icon.**

   The selected plex cannot be the last plex in the volume.

3. **From the Advanced-Ops menu, select Plex, then Detach Plexes.**

   The plex icon changes color or pattern to indicate that it is now detached. FIGURE 7-2 illustrates a detached plex (`pl-02`).

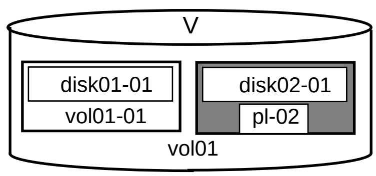

FIGURE 7-2 Detached Plex

To detach multiple plexes, select multiple plexes with the MIDDLE button and repeat the previous procedure.

7.5 Attaching a Plex

A plex that has been detached (either automatically by Volume Manager or by a user) can be attached to its parent volume again. This involves copying data from an active plex on the volume to the attaching plex. Depending on the amount of data to be copied, this operation may take some time.

One or more plexes can be attached to a parent volume at a time. In order for a plex to be attached, it must already be detached and still associated with an enabled volume.

Attach a plex as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the (detached) plex to be attached by clicking the LEFT button on its icon.**

   The selected plex must be detached.

3. **From the Advanced-Ops menu, select Plex, then Attach Plexes.**

   When the attach process completes, the attached plex icon loses the color or shading that previously indicated that it was detached.

7.6 Removing a Plex

A plex may be removed if it is no longer required for mirroring or if the space that it occupies needs to be reused. A plex can either be removed recursively (in which case any associated subdisks are automatically removed) or manually (in which case its subdisks are left behind).

Note – A plex that is associated with a volume cannot be removed until it has been dissociated from its volume.

The Basic-Ops menu provides an option that dissociates a plex if it is still associated with a volume, then recursively removes that plex and its subdisks.

When a plex is removed via the Advanced-Ops menu, its subdisk components are not automatically removed; it is necessary to manually remove these subdisks once the plex is gone. It is generally not a good idea to leave the subdisks of a removed plex behind, as the space occupied by these subdisks cannot be used by the Volume Manager's free space pool until they are removed.

Remove a plex recursively as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the plex to be removed by clicking the LEFT button on its icon.**

   The selected plex can be either associated or dissociated.
3. **From the Basic-Ops menu, select Volume Operations, then Remove Mirror.**

   If the plex is associated, it is automatically dissociated. The plex icon then is dismissed, along with any associated subdisk icons.

Remove a plex (but not its subdisks) as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **If the plex to be removed is still associated with its volume, dissociate the plex as described in the section on dissociating plexes.**
3. **Select the plex to be removed by clicking the LEFT button on its icon.**
4. **From the Advanced-Ops menu, select Plex, then Remove Plexes.**

The plex icon momentarily shrinks and is displayed next to its previous position, then is dismissed along with the subdisk icon. The plex is now removed, but its subdisk remains on its disk and is now categorized as a free subdisk.

To remove multiple plexes simultaneously, follow the previous steps, but use the MIDDLE button to select multiple plexes for removal at once.

7.7 Displaying and Altering Properties

The properties of the plex corresponding to a plex icon can be displayed (and potentially adjusted) in a properties form. The Plex Properties form provides detailed information about the characteristics of a particular plex.

To access the properties form for a particular plex icon:

- **Move the pointer onto that icon and click the RIGHT button. (If the plex is under analysis, press Shift-RIGHT instead.)**

The Plex Properties form is displayed and displays details of the plex's configuration. For detailed information on this form, refer to TABLE 7-4.

Some of the displayed properties (such as the plex's name) can be altered directly through this form by editing the appropriate field(s) and then selecting the form's Apply button. Selecting the Reset button rereads the plex properties and displays the current values.

7.8 Quick Reference for Plex Operations Menus and Forms

This section provides information on menus and forms relating to plex (mirror) operations.

7.8.1 Plex Menus

Plex operations are accessed via the Advanced-Ops menu. Most menus provide a Help selection, which contains information relevant to the items and operations listed in that particular menu.

TABLE 7-1

Menu/Submenu Access	Description
Advanced-Ops ➤ Plex	Provides access to plex (mirror) operations.
Advanced-Ops ➤ Plex ➤ Create	Creates a plex. The user may select one or more subdisks to be associated with the new plex after creation. See Also: * Plex Create Form (TABLE 7-2)
Advanced-Ops ➤ Plex ➤ Remove Plex	Removes the selected plex(es). 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.
Advanced-Ops ➤ Plex ➤ Associate Plexes	Associates one or more selected plexes with the selected volume. If the volume is started, the Volume Manager will begin to bring the plex up to date by copying all necessary data to the plex. This may take a fair amount of time. Requirements: * A volume icon and at least one plex icon must be selected. * Only non-associated plexes can be associated.
Advanced-Ops ➤ Plex ➤ Dissociate Plexes	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.

TABLE 7-1 *(Continued)*

Menu/Submenu Access	Description
Advanced-Ops ➤ Plex ➤ Attach Plexes	Attaches one or more selected plexes to their parent volumes. A plex must be detached or disabled but still associated with an enabled volume 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.
Advanced-Ops ➤ Plex ➤ Detach Plexes	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.
Advanced-Ops ➤ Plex ➤ Help	Accesses a Help window that displays information relevant to the plex operations.

7.8.2 Plex Forms

Some plex operations result in the display 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.

Some form fields are required, while others are not. Fields in these forms are read / write unless listed as read only.

7.8.2.1 Plex Create Form

This form creates a plex.

TABLE 7-2

Field	Description
Plex name:	Type 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:	Indicate the plex utility state. This is reserved for use by usage types. This field is optional.
Volume:	Type 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 14 characters. This field is optional.
Layout:	Select 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 scatters data evenly across each of its associated subdisks. The default is `Concatenated`.
Stripe unit size:	Type the unit size of the stripes on the (striped) plex. The stripe unit size 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 Striped plex layout has been specified, this field is required. This field must be blank if Concatenated plex layout has been specified.
Number of columns:	Type the number of columns to stripe across this (striped) plex. This must be a number greater than 0. This field is only applicable if a Striped or RAID-5 plex layout has been specified
Subdisks:	Displays the number of subdisks associated with the plex. This field is read only and cannot be changed.
Comment:	Allows you to add 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 Volume Manager error policies. The default is `Participate`.
Putil0:	Lists the permanent utility field 0. This is reserved for Volume Manager use, but may be changed. The maximum length of all Putil fields is 14 characters. This field is optional.
Putil1:	Lists the permanent utility field 1. This field is reserved, but may be changed. This field is optional.
Putil2:	Lists the permanent utility field 2. This field is reserved, but may be changed. This field is optional.

7.8.2.2 Stripe Gap Space Form

This form provides information about a specific region of a striped plex that contains free space. This sort of free space is visually represented either as a horizontal gap between subdisks in the same stripe column of a plex or as an indentation at the beginning of a stripe column.

A gap results in a striped plex when subdisks are removed for some reason. A subdisk equal to or smaller than the size of the gap can be dropped into the gap and thus added to the column.

Click the RIGHT mouse button on a gap between subdisk icons in the *same* column of a striped plex icon.

Note – All fields in this form are read only and cannot be changed.

TABLE 7-3

Field	Description
Plex:	The name of the striped plex containing the gap.
Column number:	The column number of the column containing the gap.
Gap offset:	The offset into the stripe column where the gap begins.
Gap size:	The size of this gap.

7.8.2.3 Plex Properties Form

This form provides detailed information on the attributes of a particular plex.

Click the RIGHT mouse button on desired plex icon. If plex is undergoing analysis, click Shift-RIGHT instead.

Plex properties can be changed via this form by altering the current values in the appropriate read / write fields and then clicking on the Apply button.

TABLE 7-4

Field	Description
Plex name:	Either use the default or type a different name for the plex. The name must be unique within this disk group. The maximum length of this field is 14 characters. The plex name can be changed by typing another name in this field.
Plex state:	Indicates the plex utility state. This is reserved for use by usage types. This field is read only and cannot be changed. This field is set to ACTIVE. However, for a RAID-5 log, this field is set to LOG.
Volume:	Lists the name of the volume that this plex should be associated with. This field is read only and cannot be changed.
Layout:	Indicates if the layout of the plex is concatenated, striped, or RAID-5. A concatenated plex is a plex with associated subdisks that are both sequentially and contiguously arranged. A striped plex is a plex that scatters data evenly across each of its associated subdisks. A RAID-5 plex also spreads data evenly across each of its associated subdisks. It provides self-contained parity information within each stripe. This field is read only and cannot be changed.
Stripe unit size:	Denotes the unit size of the stripes on the plex. If `Striped` or `RAID-5` plex layout has been specified, this field indicates the stripe unit size. This field should be blank if Concatenated plex layout has been specified. This field is read only and cannot be changed.
Stripe number:	The number of stripes on the plex. 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:	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:	Allows you to add an appropriate comment for the plex. The maximum length of the comment is 40 characters.
Errors:	Indicates whether the plex participates in Volume Manager error policies. This field is read only and cannot be changed.
Putil0:	Lists the permanent utility field 0. This is reserved for Volume Manager use, but may be changed. The maximum length of all Putil fields is 14 characters.
Putil1:	Lists the permanent utility field 1. This field is reserved, but may be changed.
Putil2:	Lists the permanent utility field 2. This field is reserved, but may be changed.
Tutil0:	Lists the temporary utility field 0. This is reserved for Volume Manager use, but may be changed. The maximum length of all Tutil fields is 14 characters.
Tutil1:	Lists the temporary utility field 1. This field is reserved, but may be changed.

TABLE 7-4 *(Continued)*

Field	Description
Tutil2:	Lists the temporary utility field 2. This field is reserved, but may be changed.
Kernel State:	Denotes the accessibility of the plex. This field is read only and cannot be changed.
Length:	Represents the length of the plex. This field is read only and cannot be changed.
Number of IO failures:	Lists the number of I/O operations that have failed on this plex. This field is read only and cannot be changed.

CHAPTER 8

Subdisk Operations

The Visual Administrator performs subdisk operations by executing the appropriate Volume Manager commands. This chapter provides instructions for performing the following subdisk operations using the Visual Administrator:

The quick reference section at the end of this chapter provides information on the Visual Administrator subdisk operations menus and forms.

FIGURE 8-1 illustrates a subdisk associated with a plex in a volume (`disk01-01`) and a free subdisk (`disk02-01`).

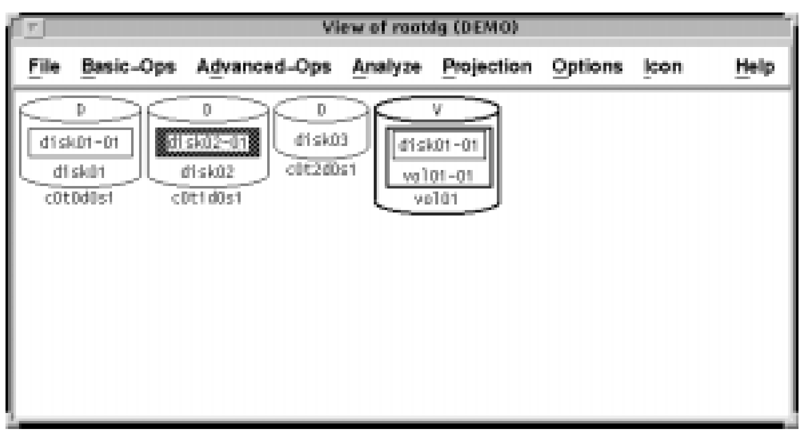

FIGURE 8-1 Subdisks

8.1 Showing Free Subdisks

The Projection menu provides access to a feature that highlights any free subdisk icons. This is useful for identifying subdisks that are not currently associated with any plexes and should either be used or removed to free up the space they occupy.

To turn on highlighting of free subdisks:

1. **Click LEFT on the Projection menu.**
2. **Click LEFT on Show Free Subdisks from the Projection menu.**
3. **Click LEFT on Start from the Show Free Subdisks menu.**

   Once turned on, free subdisks will continue to be highlighted until this feature is turned off.

8.2 Creating a Subdisk

Subdisks are the lowest-level building blocks in a volume and need to be created to build a volume. Subdisks are created on selected VM disks. Once created, subdisks can be associated with plexes to build volumes.

When building mirrored or striped volumes, care must be taken to create subdisks on separate disks in order for striping or mirroring to work properly.

Create a subdisk as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the disk on which to create the subdisk.**
3. **From the Advanced-Ops menu, select Subdisk, then Create.**

   The Subdisk Create form is displayed. This form establishes the parameters for the subdisk. Some of the fields already contain default values and others are optional; only the subdisk length is required.

4. **Complete the Subdisk Create form, typing the subdisk length and optionally altering other fields.**

   For detailed information on this form, refer to online help or TABLE 8-2.

5. **When the form is properly completed, select Apply to activate the subdisk creation.**

   A new subdisk icon is displayed in the selected disk icon.

8.3 Associating a Subdisk

Once a subdisk has been created, it can be associated with a plex, which can then be used to build a volume. It is not possible to read or write to subdisks alone.

Several subdisks can be associated with a single plex at a time, but only one plex can be selected per association operation. When a subdisk is associated with a plex that is already associated with a volume, the plex length is automatically increased by the subdisk length, but the volume length is not updated.

There are a few methods of associating a subdisk with a plex:

- Selecting the plex and subdisk(s) and then using menus.
- Dragging each subdisk icon onto the plex icon and then dropping the subdisk icon on that plex when the pointer (in the image of a hand) is positioned over an unobscured portion of the plex icon.
- Dragging a subdisk icon onto a gap (or "hole" formed by the removal of a subdisk) between two subdisks on a plex and dropping the subdisk there. If there is not sufficient space in the hole for the new subdisk, an error will result.
- Automatically associating subdisks with a plex by selecting the subdisk icon(s) immediately prior to the creation of a plex. See Chapter 7, "Plex Operations," for further details.

Associate a subdisk with an existing plex via menus as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select both the subdisk and the plex to which it is to be associated by clicking the MIDDLE button on both of their icons.**
3. **From the Advanced-Ops menu, select Subdisk, then Associate Subdisks.**

   The subdisk icon is displayed within the plex icon. The subdisk is now associated with the plex.

   To associate multiple subdisks with a single plex, select multiple subdisks along with the plex and repeat the previous procedure.

8.3.1 Log Subdisks

A log subdisk is used to log recent disk activity with Dirty Region Logging (described in Chapter 1, “Introduction to the Volume Manager”) in effect. A subdisk can be associated with a plex as a log subdisk. Only one log subdisk can exist for a given plex.

Associate a log subdisk with an existing plex as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Create a subdisk, as described previously.**
3. **Select both the subdisk and the plex to which it is to be associated by clicking the MIDDLE button on both of their icons.**
4. **From the Advanced-Ops menu, select Subdisk, then Associate as Log Sd.**

   The log subdisk icon is displayed within the plex icon and is distinguished from normal subdisk icons by its double border. The subdisk is now associated with the plex as a log subdisk.

Note – A log subdisk can also be created at the same time a volume is created. This is done by setting the Create log subdisk field in the Simple or Striped Volume/FS Create Form to `Yes`.

8.4 Dissociating a Subdisk

Associated subdisks can be dissociated from their parent plexes. A subdisk becomes a free subdisk when it is dissociated. A subdisk must be dissociated before it can be removed.

There are a couple of methods of dissociating a subdisk from a plex:

- Selecting the subdisk(s) and then using menus.
- Dragging the subdisk icon beyond the borders of the plex icon and then dropping the subdisk in an open area of the view window.

Both normal and log subdisks are dissociated in the same way.

Dissociate a subdisk from its plex via menus as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the subdisk to be dissociated from its plex by clicking the LEFT button on its icon.**

3. **From the Advanced-Ops menu, select Subdisk, then Dissociate Subdisks.**

   The subdisk icon is dismissed from its plex icon and is only visible in its VM disk icon. The subdisk is now dissociated from its plex.

   To dissociate multiple subdisks from a plex, select multiple subdisks (using the MIDDLE button) and repeat the previous procedure.

8.5 Evacuating Subdisks

The evacuation of subdisks involves moving all subdisks from one disk to another. This is useful when a disk appears to be failing and you want to protect any volumes that rely on that disk. This operation can only be performed between two disks in the same disk group.

Evacuate subdisks from a disk as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**

2. **Select the disk *from* which subdisks are to be evacuated by clicking LEFT on its icon.**

3. **From the Basic-Ops menu, select Disk Operations, then Evacuate Subdisks.**

   The Evacuate Subdisks form is displayed.

4. **Complete the Evacuate Subdisks form.**

   For detailed information on this form, refer to online help or TABLE 8-3.

5. **When the form is properly completed, select Apply to activate the subdisk evacuation.**

   The subdisks move from their original disk icon to the targeted disk icon.

8.6 Removing a Subdisk

A subdisk may be removed to free its space for reuse. Selected subdisks can be permanently removed. Once a subdisk has been dissociated from a plex, it is generally better to remove that subdisk than to allow it to continue to consume disk space. The space occupied by free subdisks cannot be used by the Volume Manager's free space pool until those subdisks are removed.

Note – A subdisk that is associated with a plex cannot be removed; the subdisk must be dissociated from its plex beforehand.

Remove a subdisk as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **If the subdisk to be removed is still associated with its plex, dissociate the subdisk as described in the section on dissociating subdisks.**
3. **Identify the subdisk icon to be removed.**

   Since the subdisk is no longer associated with a plex, its icon should now exist only on its disk. If projection is set to show free subdisks (via the Projection menu), then this subdisk icon should be highlighted on its disk.
4. **Select the subdisk to be removed by clicking the LEFT button on its icon.**
5. **From the Advanced-Ops menu, select Subdisk, then Remove Subdisks.**

   The subdisk icon is dismissed, leaving a gap or open area on its disk.

   To remove multiple subdisks simultaneously, follow the previous steps, but use the MIDDLE button to select multiple subdisks for removal at once.

8.7 Splitting, Moving, and Joining Subdisks

In some situations, it is useful to be able to split subdisks into a series of smaller pieces. This makes it possible to more accurately identify those areas of a disk where hot spots exist. Analysis can be used to identify any subdisk areas associated with high activity. The subdisks occupying the hot spot location can then be moved to another disk for improved overall performance. Subdisks that were not occupying hot spots can later be rejoined to form larger subdisks.

Split a subdisk into three parts as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select the subdisk to be split by clicking the LEFT button on its icon.**
3. **From the Advanced-Ops menu, select Subdisk, then Split the Subdisk. From the resulting submenu, select Into More Than 2 Subdisks.**

   The Subdisk Split form is displayed.
4. **Complete the Subdisk Split form, specifying `3` as the number of new subdisks.**

   For detailed information on this form, refer to online help or TABLE 8-5.
5. **When the form is properly completed, select Apply to activate the subdisk split.**

   The subdisk icon is converted into 3 subdisk icons.

   Analysis can now be performed on the resulting three subdisks (using the Analyze menu) to determine which of those subdisks is the hottest spot. Once a subdisk is identified as a hot spot, it can be moved to another disk.

Move a subdisk to another VM disk as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Drag the desired subdisk icon from one VM disk to a second disk and drop it there.**

   A free subdisk identical to the original subdisk is created on the new disk and contents of the original subdisk are copied to the free subdisk. As the original subdisk is copied to its new disk, a temporary plex is displayed. When the transfer is complete, the original subdisk is removed.

Contiguous subdisks on any disk can be joined together. Joining does not succeed for non-contiguous subdisks.

Join a pair of contiguous subdisks as follows:

1. **Go to the view window corresponding to the disk group in which you want to perform this operation (View of `rootdg`, by default).**
2. **Select both subdisks by clicking the MIDDLE button on their icons within their disk icon.**

   The subdisks must be adjacent on a single disk, and if associated, must be to the same plex in order for joining to occur.

3. **From the Advanced-Ops menu, select Subdisk, then Join subdisks.**

   The subdisk icons rewrite themselves on their disk and reappear as a single subdisk icon. The subdisks are now joined.

8.8 Displaying and Altering Properties

The properties of the subdisk corresponding to a subdisk icon can be displayed (and potentially altered) in a properties form. The Subdisk Properties form provides detailed information about the characteristics of a particular subdisk.

To access the properties form for a particular subdisk icon:

- **Move the pointer onto that icon and click the RIGHT button. (If the subdisk is under analysis, press Shift-RIGHT instead.)**

  The Subdisk Properties form displays details of the subdisk's configuration. For detailed information on this form, refer to online help or TABLE 8-6.

Some of the displayed properties (such as the subdisk's name) can be altered directly through this form by editing the appropriate field(s) and then selecting the form's Apply button. Selecting the Reset button rereads the subdisk properties and displays the current values.

8.9 Quick Reference to Subdisk Operations Menus and Forms

This section provides information menus and forms relevant to subdisk operations.

8.9.1 Subdisk Menus

The Advanced-Ops menu provides access to subdisk-related menus. Most menus provide a Help selection, which contains information relevant to the items and operations listed in that particular menu.

TABLE 8-1

Menu/Submenu Access	Description
Advanced-Ops ➤ Subdisk	Provides access to subdisk operations.
Advanced-Ops ➤ Subdisk ➤ Create	Creates a subdisk on the selected VM disk. Requirements: * A VM disk must be selected. See Also: * Subdisk Create Form (TABLE 8-2)
Advanced-Ops ➤ Subdisk ➤ Remove Subdisk	Removes the selected subdisk(s). 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.
Advanced-Ops ➤ Subdisk ➤ Associate Subdisks	Associates one or more subdisks with the selected plex (mirror). Requirements: * A plex icon and at least one subdisk icon must be selected. * Only non-associated (free) subdisks can be associated.
Advanced-Ops ➤ Subdisk ➤ Associate as Log Sd	Associates the selected subdisk as a log subdisk with the selected plex (mirror). Dirty Region Logging is in effect. 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 non-associated (free) subdisks can be associated. * The selected plex cannot already have a log subdisk.
Advanced-Ops ➤ Subdisk ➤ Dissociate Subdisks	Dissociates one or more selected subdisks from their parent plexes (mirrors). 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.

TABLE 8-1 *(Continued)*

Menu/Submenu Access	Description
Advanced-Ops ➤ Subdisk ➤ Join Subdisks	Joins the selected subdisks together 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.
Advanced-Ops ➤ Subdisk ➤ Split the Subdisk	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 (mirror), the resulting subdisks will also be associated with that plex. From the Split the Subdisk menu, a cascading menu enables you to indicate whether the subdisk is to be split into two or several parts. Requirements: * Only one subdisk icon can be selected. See Also: * Subdisk Split Into Two (TABLE 8-4) * Subdisk Split Into Many (TABLE 8-5)
Advanced-Ops ➤ Subdisk ➤ Split the Subdisk ➤ Into 2 Subdisks	Splits the selected subdisk into 2 subdisks.
Advanced-Ops ➤ Subdisk ➤ Split the Subdisk ➤ Into More Than 2 Subdisks	Splits the selected subdisk into several subdisks. Requirements: * The number of subdisks is limited by the amount of space left in the configuration database.
Advanced-Ops ➤ Subdisk Help	Accesses a Help window, which displays information relevant to the subdisk operations.

8.9.2 Subdisk Forms

Some subdisk operations result in the display 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.

Some form fields are required, while others are not. Fields on these forms are read/write unless listed as read only.

8.9.2.1 Subdisk Create Form

This form creates a subdisk according to the your specifications.

TABLE 8-2

Field	Description
Disk name:	Displays the name of the Volume Manager disk on which the subdisk is to be created. This field is read only.
Subdisk name:	Either use the default or type a different name for 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:	Type 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:	Type 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:	Type the name of the plex (mirror) 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:	Indicate the offset of this subdisk into its associated plex (mirror). 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.
Plex column:	If the named plex is striped, this field indicates the column number (0 or greater) to which this subdisk is to be assigned. This field is ignored if no plex is specified or the plex is not a striped plex.
Comment:	Enables you to add 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 Volume 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.

8.9.2.2 Evacuate Subdisk Form

This form is used to evacuate subdisks. This operation can only be performed between two disks in the same disk group.

TABLE 8-3

Field	Description
Disk group name:	Type the name of the disk group to which both disks belong. Both disks must share the same disk group.
Evacuate From:	Displays the name of the disk from which the subdisks are to be evacuated, which can be changed.
To:	Type the name of the disk to which the subdisks are to be moved. This field is optional. However, the evacuated subdisks will be moved to one or more random disks if no target disk is specified.

8.9.2.3 Subdisk Split Into Two Form

This form is used to split the selected subdisk into exactly 2 subdisks. The first subdisk retains the name of the original one; the second subdisk adopts the name specified in this form.

TABLE 8-4

Field	Description
Present size:	Displays the size of the subdisk to be split. This field is read only and cannot be changed.
Name of new subdisk:	Type 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:	Type 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.

8.9.2.4 Subdisk Split Into Many Form

This form is used to split the selected subdisk into several subdisks of equal sizes. The first subdisk retains the name of the original one; the additional subdisks are automatically named by the Volume Manager.

TABLE 8-5

Field	Description
Original subdisk:	Displays the name of the selected subdisk. This field is read only.
Present size:	Displays 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.
Number of new subdisks:	Indicate 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.

8.9.2.5 Subdisk Properties Forms

This form provides detailed information on the attributes of a particular subdisk. 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 the Apply button.

Click the RIGHT mouse button on desired subdisk icon. (If the subdisk is undergoing analysis, press Shift-RIGHT instead.)

TABLE 8-6

Field	Description
Disk name:	Displays the name of the disk where the subdisk resides. This field is read only.
Subdisk name:	Displays 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 typing another name in this field.
Disk offset:	Indicates the length into the disk where this subdisk is located, in sectors. This field is read only.
Subdisk length:	Specify 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:	Displays the name of the plex (mirror) with which the subdisk is associated. This field is read only.

TABLE 8-6 *(Continued)*

Field	Description
Plex offset:	Displays the offset of this subdisk into its associated plex (mirror). If the subdisk is not associated, this field contains a zero. This field is read only and cannot be changed.
Column number:	Specify the number of the stripe column in which this subdisk is located. This field applies only to subdisks that are part of a striped or a RAID-5 plex.
Comment:	Enables you to add 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:	Lists the permanent utility field 0. This is reserved for Volume Manager use, but may be changed. The maximum length of all Putil fields is 14 characters.
Putil1:	Lists the permanent utility field 1. This field is reserved, but may be changed. The maximum length of this field is 14 characters.
Putil2:	Lists the permanent utility field 2. This field is reserved, but may be changed. The maximum length of this field is 14 characters.
Tutil0:	Lists the temporary utility field 0. This is reserved for Volume Manager use, but may be changed. The maximum length of all Tutil fields is 14 characters.
Tutil1:	Lists the temporary utility field 1. This field is reserved, but may be changed. The maximum length of this field is 14 characters.
Tutil2:	Lists the 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:	Lists the number of I/O operations that have failed on this subdisk. This field is read only and cannot be changed.

CHAPTER **9**

CLI Introduction

This chapter provides some preliminary information that you should know before attempting to use the Volume Manager Command Line Interface (CLI).

Refer to the chapters that follow for information on using Volume Manager commands to perform common tasks. The `vxintro`(1M) manual page also contains introductory information relating to Volume Manager utilities.

Topics covered in this chapter include

9.1 Disk Naming

When performing disk administration, it is important that you recognize the difference between a *device name* and a *disk name.*

The *device name* (sometimes referred to as *devname* or *disk access name*) is the location of the disk. The syntax of a device name is `c#t#d#s#`, where:

- `c#` is the number of the controller to which the disk drive is attached.
- `t#` is the number of the target disk on that controller.
- `d#` is the number of the disk (or UNIX partition).
- `s#` is the number of the disk slice.

The full path name of a device is `/dev/vx/dmp/`*devicename*. In this document, only the device name is listed and `/dev/vx/dmp` is assumed. An example of a device name is `c0t0d0s2`.

The *disk name* (sometimes referred to as *disk media name*) is an administrative name for the disk, such as `disk01`. If you do not assign a disk name, the disk name defaults to `disk##` if the disk is being added to `rootdg` (where ## is a sequence number). Otherwise, the default disk name is *groupname##*, where *groupname* is the name of the disk group to which the disk is added.

9.2 Displaying Configuration Information

9.2.1 Displaying Disk Listings

To see a list of available disks, use the following command:

```
vxdisk list
```

This command provides a list of all disks available to the Volume Manager. The output of `vxdisk list` lists the device name, the type of disk, the disk name, the disk group to which the disk belongs, and the disk's status. The following is an example of `vxdisk list` output:

```
DEVICE      TYPE      DISK      GROUP      STATUS
c0t0d0s2    sliced    disk01    rootdg     online
c0t1d0s2    sliced    disk02    rootdg     online
c0t2d0s2    sliced    disk03    rootdg     online
```

9.2.2 Displaying VxVM Object Listings

The vxprint command displays detailed information on existing Volume Manager objects. Display detailed output for all currently-existing objects as follows:

```
vxprint -ht
```

This displays output such as the following:

```
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       774226267.1025.tweety

dm disk01       c0t0d0s2     sliced   1055     1054944  -
dm disk02       c0t1d0s2     sliced   1119     1043840  -

v  vol03        fsgen        DISABLED ACTIVE   1310720  SELECT    -
pl vol03-01     vol03        DISABLED ACTIVE   1311632  STRIPE    2/128    RW
sd disk01-01    vol03-01     disk01   0        655380   0/0       c0t0d0   ENA
sd disk02-08    vol03-01     disk02   0        655760   1/0       c0t1d0   ENA

v  vm-build     fsgen        DISABLED ACTIVE   409600   SELECT    -
pl vm-build-01  vm-build     DISABLED ACTIVE   410016   STRIPE    2/128    RW
sd disk01-02    vm-build-01  disk01   655380   204831   0/0       c0t0d0   ENA
sd disk02-01    vm-build-01  disk02   655760   204960   1/0       c0t1d0   ENA
```

where `dg` is a disk group, `dm` is a disk, `v` is a volume, `pl` is a plex, and `sd` is a subdisk. The top few lines indicate the headers that match each type of output line that follows. Each volume is listed along with its associated plex(es) and subdisk(s).

9.2.3 Displaying Free Space in a Disk Group

Before you add volumes and file systems to your system, you may want to make sure you have enough free disk space to adequately meet your needs. The Volume Manager lets you request a display of free space.

To display free space for a disk group, type:

```
vxdg -g diskgroup free
```

where `-g` *diskgroup* optionally specifies a disk group.

To see the free space in the default disk group, `rootdg`, type:

```
vxdg free
```

The Volume Manager returns the following:

```
GROUP     DISK     DEVICE      TAG       OFFSET    LENGTH    FLAGS
rootdg    disk04   c0t0d0s2    c0t0d0    726400    102672    -
rootdg    disk01   c1t0d0s2    c1t0d0    0         102128    -
rootdg    disk01   c1t0d0s2    c1t0d0    175856    26384     -
rootdg    disk02   c1t1d0s2    c1t1d0    26624     175616    -
```

The free space is measured in 512-byte sectors.

9.3 Creating Volumes

Volumes need to be created to take advantage of the Volume Manager concept of virtual disks. Once a volume exists, a file system can be placed on the volume to organize the disk space with files and directories. Alternatively, applications like databases might be used to organize data on volumes.

You can create volumes using either a basic or advanced approach:

- Basic — The basic approach takes information about what the user wants to accomplish and then performs the necessary underlying tasks. This approach requires only minimal input from the user, but does permit more detailed specifications. Basic operations are performed primarily through the `vxassist` command.
- Advanced — The advanced approach consists of a number of fairly complicated commands that typically require the user to specify detailed input. These commands use a “building block” approach that requires the user to have a detailed knowledge of the underlying structure and components to manually perform the sequences of commands necessary to accomplish a certain task. Advanced operations are performed through several VxVM commands.

The creation of a volume involves the creation of plex and subdisk components. With the basic approach to volume creation, the underlying plexes and subdisks are handled automatically. You simply indicate the desired volume characteristics.

Volumes can be created with various layout types:

- Simple — A concatenated volume whose subdisks are arranged both sequentially and contiguously within a plex. Concatenation enables a volume to be created from multiple regions of one or more disks if there is not enough space for an entire volume on a single region of a disk.
- Striped — A volume with data spread fairly evenly across multiple disks. *Stripes* are equally-sized fragments that are allocated alternately and evenly to the subdisks of a single plex. There should be at least two subdisks in a striped plex, each of which should exist on a different disk. Throughput increases with the number of disks across which a plex is striped. Striping helps to balance I/O load in cases where high traffic areas exist on certain subdisks.
- RAID-5 — A volume that uses striping to spread data fairly evenly across multiple disks in an array, and enables independent access. It also stripes parity across all the disks in the array, rather than being concentrated on a single parity disk. Each stripe contains a parity stripe unit and data stripe units. In comparison to the performance of striped volumes, write throughput of RAID-5 volumes decreases since parity information needs to be updated each time data is accessed. However, in comparison to mirroring, the use of parity reduces the amount of space required.
- Mirrored — A volume with multiple plexes that duplicate the information contained in a volume. Although a volume can have a single plex, at least two are required for true mirroring (redundancy of data) to be in effect. Each of these plexes should contain disk space from different disks, in order for the redundancy to be useful.
- Striped and Mirrored — A volume with a striped plex and another plex that mirrors the striped one. This requires at least two disks for striping and one or more other disks for mirroring (depending on whether the plex is simple or striped). A striped and mirrored volume is advantageous because it both spreads data across multiple disks and provides redundancy of data.

9.4 General Use of the `vxassist` Utility

The `vxassist` utility provides a simple way to create and manage volumes from the command line. You can use `vxassist` to create, mirror, grow, shrink, and back up volumes. `vxassist` is capable of performing many operations that would otherwise require the use of a series of more complicated VxVM commands. `vxassist` creates and manipulates volumes based on a set of established defaults, but also enables the user to supply preferences for each operation.

The `vxassist` command typically takes the following form:

```
vxassist keyword volume_name [attributes...]
```

Select the specific action to perform by specifying an operation keyword as the first argument on the command line. For example, the keyword for creating a new volume is `make`.

To create a new volume, type:

```
vxassist make volume_name length
```

The first argument after any `vxassist` keyword is always a volume name. Follow the volume name with a set of attributes. Use these attributes to specify where to allocate space and whether you want mirroring or striping to be used.

You can select the disks on which the volumes are to be created by specifying the disk names at the end of the command line. For example, to create a 30 megabyte striped volume on three specific disks (`disk03`, `disk04`, and `disk05`), type:

```
vxassist make stripevol 30m layout=stripe disk03 disk04 disk05
```

`vxassist` defaults are listed in the file `/etc/defaults/vxassist`. The defaults listed in this file take effect if there is no overriding default specified on the `vxassist` command line.

Refer to the *Sun StorEdge Volume Manager 2.6 System Administrator's Guide* or the `vxassist`(1M) manual page for more information on operations, options, and attributes of the `vxassist` utility.

CHAPTER **10**

Disk CLI Operations

This chapter provides instructions on performing the following disk operations using the command line interface:

Refer to the *Sun StorEdge Volume Manager 2.6 System Administrator's Guide* for additional information and examples on using the command line interface to perform disk operations. For more information about physical disks, VM disks, and disk groups, see Chapter 1, "Introduction to the Volume Manager."

10.1 Physical Disk Operations

Physical disk operations are described in the following sections.

10.1.1 Bringing a Physical Disk Under Volume Manager Control

When you add a disk to a system that is running the Volume Manager, you may want to put the disk under control of the Volume Manager so that it can control the space allocation on the disk. If the disk was previously in use, but not under Volume Manager control, then you may want to preserve existing data on the disk while still letting the Volume Manager take control of the disk. This can be accomplished using the encapsulation function of the Volume Manager. If the disk is new, then it will need to be initialized. If a disk was previously not under Volume Manager control, but no data is required to be preserved, an initialization operation should also be performed.

Add a disk by typing the command:

```
vxdiskadd devname
```

where *devname* is the device name of the disk to be added.

To add the device `c1t0d0` to Volume Manager control, do the following:

1. **Type the following to start** `vxdiskadd`**:**

```
vxdiskadd c1t0d0
```

Notice that the `s2` suffix is not used here.

2. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
Add or initialize disks
Menu: VolumeManager/Disk/AddDisks

 Here is the disk selected.  Output format: [Device_Name]

  c1t0d0

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

3. **At the following prompt, specify the disk group to which the disk should be added or press Return to accept** `rootdg`**:**

```
  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)
```

4. **At the following prompt, either press Return to accept the default disk name or type a disk name:**

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

5. **When prompted whether this disk should become a hot-relocation spare, type** `n` **(or press Return):**

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

6. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
  The selected disks will be added to the disk group rootdg with
  default disk names.

  c1t0d0

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

7. **If there is data on this disk that needs to be preserved, type** `y` **to select encapsulation:**

```
  The following disk device has a valid VTOC, but does not appear to have
  been initialized for the Volume 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.
  Output format: [Device_Name]

  c1t0d0

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

8. **To continue the operation, press Return at the following prompt:**

```
  The following disk has been selected for encapsulation.
  Output format: [Device_Name  c1t0d0]

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

A message similar to the following indicates that the disk is being encapsulated for VxVM use:

```
   The disk device c1t0d0 will be encapsulated and added to the disk group
  rootdg with the disk name disk01.

  The c1t0d0 disk has been configured for encapsulation.

  The first stage of encapsulation has completed successfully.  You
  should now reboot your system at the earliest possible opportunity.

  The encapsulation will require two or three reboots which will happen
  automatically after the next reboot.  To reboot execute the command:

shutdown -g0 -y -i6

  This will update the /etc/vfstab file so that volume devices are
  used to mount the file systems on this disk device.  You will need
  to update any other references such as backup scripts, databases,
  or manually created swap devices.

Goodbye.
```

Shutdown and reboot your system as soon as convenient.

10.1.2 Reserving Physical Disks

By default, `vxassist` allocates space from any disk that has free space. You may want to reserve some set of disks for special purposes, such as to avoid general use of a particularly slow or a particularly fast disk.

To reserve a disk for special purposes, type:

```
vxedit set reserve=on diskname
```

After you type this command, `vxassist` will not allocate space from the selected disk unless that disk is specifically mentioned on the `vxassist` command line. For example, if disk `disk03` is reserved, the command:

```
vxassist make vol03 20m disk03
```

overrides the reservation and creates a 20 megabyte volume on `disk03`. However, the command:

```
vxassist make vol04 20m
```

does not use `disk03`, even if there is no free space on any other disk.

To turn off reservation of a disk, type:

```
vxedit set reserve=no diskname
```

10.1.3 Adding a Physical Disk to a Disk Group

You may want to add a new disk to an already established disk group. Perhaps the current disks have insufficient space for the application or work group requirements, especially if these requirements have changed.

To add an initialized disk to a disk group, type:

```
vxdiskadd devname
```

To add device `c1t1d0` to `rootdg`, do the following:

1. **Type the following to start** `vxdiskadd`**:**

```
vxdiskadd c1t1d0
```

vxdiskadd displays the following message:

```
Add or initialize disks
Menu: VolumeManager/Disk/AddDisks

 Here is the disk selected.  Output format: [Device_Name]

  c1t1d0

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

2. **At the following prompt, specify the disk group to which the disk should be added or press Return to accept** `rootdg`**:**

```
  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)
```

3. **At the following prompt, either press Return to accept the default disk name or type a disk name:**

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

4. **When** `vxdiskadd` **asks whether this disk should become a hot-relocation spare, type** `n` **(or press Return):**

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

5. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
  The selected disks will be added to the disk group rootdg with
  default disk names.

  c1t1d0

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

6. **The following prompt indicates that this disk has been previously initialized for future VxVM use; type `y` to confirm that you now want to use this disk:**

```
  The following disk device appears to have been initialized already.
  The disk is currently available as a replacement disk.
  Output format: [Device_Name]

  c1t1d0

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

7. **To reinitialize the disk, type `y` (or press Return) at the following prompt:**

```
  The following disk you selected for use appears to already have
  been initialized for the Volume Manager.  If you are certain the
  disk has already been initialized for the Volume Manager, then you
  do not need to reinitialize the disk device.
  Output format: [Device_Name]

  c1t1d0

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

Messages similar to the following should now confirm that this disk is being reinitialized for VxVM use.

```
  Initializing device c1t1d0.

Adding disk device c1t1d0 to disk group rootdg with disk
  name disk03.
```

To confirm that the disk has been added to the disk group, type:

```
vxdisk list
```

The Volume Manager returns a listing similar to the following:

```
DEVICE       TYPE     DISK      GROUP     STATUS
c0t0d0s2     sliced   disk04    rootdg    online
c1t0d0s2     sliced   disk01    rootdg    online
c1t1d0s2     sliced   disk03    rootdg    online
```

10.1.4 Taking a Physical Disk Offline

Occasionally, you may need to take a physical disk offline. If the disk is corrupted, you need to disable it and remove it. You also must disable a disk before moving the physical disk device to another location to be connected to another system.

To take a physical disk offline, first remove the disk from its disk group. Then place the disk in an ‘‘offline’’ state as follows:

```
vxdisk offline devname
```

To take the device `c1t1d0s2` off line, type:

```
vxdisk offline c1t1d0s2
```

Note – The device name is used here because the disk is no longer in a disk group and so does not have an administrative name.

10.1.5 Removing a Physical Disk

You can remove a disk to move it to another system or you may remove the disk because the disk is failing or has failed. However, before removing the disk from the current system, you must:

1. Unmount any file systems on the volumes.
2. Stop the volumes.
3. Move the volumes to other disks or back up the volumes to tape. To move a volume, mirror the volume on one or more other disks, then remove the original copy of the volume.

Alternatively, if the volumes are no longer needed, they can be removed.

Removing a disk involves the following steps:

1. **Remove a disk from a disk group as follows:**

```
vxdg [-g groupname] rmdisk diskname
```

where *groupname* is the name of the group to which the disk belongs and *diskname* is the name of the disk to be removed.

For example, to remove `disk01` from `rootdg`, type:

```
vxdg rmdisk disk01
```

Since `rootdg` is the default disk group, you do not need to specify it.

2. **Remove the disk from Volume Manager control as follows:**

```
vxdisk rm devname
```

For example, to remove `c1t0d0` from Volume Manager control, type:

```
vxdisk rm c1t0d0s2
```

10.2 VM Disk Operations

VM disk operations are described in the following sections.

10.2.1 Displaying Disk Information

Before you use a disk, you need to know if it has been initialized and placed under Volume Manager control. You also need to know if the disk is part of a disk group, since you cannot create volumes on a disk that is not part of a disk group. The `vxdisk list` command displays device names for all recognized disks, the disk names, the disk group names associated with each disk, and the status of each disk.

To display information on all disks that are defined to the Volume Manager, type:

```
vxdisk list
```

The Volume Manager returns the following display:

```
DEVICE      TYPE      DISK       GROUP       STATUS
c0t0d0s2    sliced    disk04     rootdg      online
c1t0d0s2    sliced    disk01     rootdg      online
c1t1d0s2    sliced    -          -           online
```

10.2.2 Displaying Multipaths Under a VM Disk

The `vxdisk` utility is used to display the multipathing information for a particular metadevice. The metadevice is a device representation of a particular physical disk having mutiple physical paths from the I/O controller of the system. In the Volume Manager all the physical disks in the system are represented as metadevices with one or more physical paths.

You can see the multipathing information for a particular metadevice with the command:

```
vxdisk list c1t5d0s2
```

`c1t5d0s2` is the metadevice formed by the Volume Manager.

10.2.3 Displaying Multipathing Information

To display details on a particular disk defined to the Volume Manager, type:

```
vxdisk list disk01
```

The Volume Manager returns the following display:

```
Device    c2t0d0s2
devicetag c2t0d0
type      sliced
hostid    aparajita
disk      name=disk01 id=861086917.1052.aparajita
group     name=rootdg id=861086912.1025.aparajita
flags     online ready autoconfig autoimport imported
pubpaths  block=/dev/vx/dmp/c2t0d0s4 char=/dev/vx/rdmp/c2t0d0s4
privpaths block=/dev/vx/dmp/c2t0d0s3 char=/dev/vx/rdmp/c2t0d0s3
version   2.1
iosize    min=512 (bytes) max=2048 (blocks)
public    slice=4 offset=0 len=1043840
private   slice=3 offset=1 len=1119
update    time=861801175 seqno=0.48
headers   0 248
configs   count=1 len=795
logs      count=1 len=120
Defined regions
config   priv 000017-000247[000231]:copy=01 offset=000000 enabled
config   priv 000249-000812[000564]:copy=01 offset=000231 enabled
log     priv 000813-000932[000120]:copy=01 offset=000000 enabled
Multipathing information:
numpaths:       2
c2t0d0s2        active
c1 t0d0s2        failed
```

The output shows two paths to a physical device represented by the metadevice `c2t0d0s2`. The path `c2t0d0s2` is active and the other path `c1t0d0s2` is in a failed state.

10.2.4 Adding a VM Disk to the Hot-Relocation Pool

Hot-relocation is the ability of a system to automatically react to I/O failure by relocating redundant subdisks to other disks and restoring the affected VxVM objects and data. If a disk has already been designated as a spare in the disk group, the subdisks from the failed disk are relocated to the spare disk. Otherwise, any suitable free space in the disk group is used. Refer to Chapter 1, “Introduction to the Volume Manager” for more information.

To designate a disk as a hot-relocation spare, type:

```
vxedit set spare=on diskname
```

For example, to designate disk `disk01` as a spare, type:

```
vxedit set spare=on disk01
```

You can use the `vxdisk list` command to confirm that this disk is now a spare; `disk01` should be listed with a `spare` flag.

Any VM disk in this disk group can now use this disk as a spare in the event of a failure. If a disk fails, hot-relocation should automatically occur (if possible). You should be notified of the failure and relocation via electronic mail. After successful relocation, you may want to replace the failed disk.

10.2.5 Removing a VM Disk From the Hot-Relocation Pool

While a disk is designated as a spare, the space on that disk is not used as free space for the creation of VxVM objects within its disk group. If necessary, you can free a spare disk for general use by removing it from the pool of hot-relocation disks.

To determine which disks are currently designated as spares, use the command `vxdisk list`. The output of this command should list any spare disks with the `spare` flag.

To remove a spare from the hot-relocation pool, type:

```
vxedit set spare=off diskname
```

For example, to make disk `disk01` available for normal use, type:

```
vxedit set spare=off disk01
```

10.2.6 Renaming a VM Disk

It is not necessary to give your disks special names. The Volume Manager gives the disk a default name when you add the disk to Volume Manager control. The disk name is used by the Volume Manager to identify the disk's location or type. If you want to change the disk name to reflect a change of ownership or use, type:

```
vxedit rename old_diskname new_diskname
```

To rename `disk01` to `disk03`, type:

```
vxedit rename disk01 disk03
```

To see if the name change took place, type:

```
vxdisk list
```

The Volume Manager returns the following:

```
DEVICE    TYPE    DISK     GROUP    STATUS
c0t0d0s2  sliced  disk04   rootdg   online
c1t0d0s2  sliced  disk03   rootdg   online
c1t1d0s2  sliced  -        -        online
```

Note – By default, VxVM names subdisk objects after the VM disk on which they are located. Renaming a VM disk does not automatically rename the subdisks on that disk.

10.3 Mirroring the Boot Disk

To mirror your boot (root) disk onto another disk, use either the steps outlined here or `vxdiskadm` (see Chapter 13, "Menu Interface Operations"). This makes it possible to recover from failure of your boot disk by replacing it with the mirror of the boot disk.

Before mirroring your boot disk, the EEPROM variable `use-nvramrc?` must be set to `true` if you want to take advantage of the Volume Manager boot disk aliases to identify the mirror of the boot disk if a replacement is needed. If this variable is set to `false`, you will need to determine which disks are bootable yourself. You can set this variable to `true` as follows:

```
eeprom use-nvramrc?=true
```

To mirror your boot disk, do the following:

1. **Select a disk that is at least as large as your boot disk.**
2. **Use the** `vxdiskadd` **command to add the selected disk as a new disk (if it is not already added).**
3. **Execute the following command:**

```
/etc/vx/bin/vxrootmir alternate_disk
```

where *alternate_disk* is the disk name assigned to the disk.

vxrootmir creates a mirror for rootvol (the volume for the root file system).

The alternate boot disk is configured to enable booting from it if the primary boot disk fails.

There may be other volumes on the boot disk, such as volumes for /home or /tmp file systems. These can be mirrored separately using the vxassist utility. For example, if you have a /home file system on a volume homevol, you can mirror it to *alternate_disk* using the command:

```
vxassist mirror homevol alternate_disk
```

If you do not have space for a copy of some of these file systems on your alternate boot disk, you can mirror them to other disks. You can also span or stripe these other volumes across other disks attached to your system.

To list all volumes on your primary boot disk, use the command:

```
vxprint -t -v -e 'aslist.aslist.sd_disk="boot_disk_name"'
```

To mirror all of the simple volumes on this disk to your alternate boot disk, use the command:

```
/etc/vx/bin/vxmirror boot_disk_name alternate_disk
```

10.4 Disk Group Operations

Disk group operations are described in the following sections.

10.4.1 Initializing a New Disk Group

There can be situations in which all data related to a particular set of applications or a particular group of users needs to be made accessible on another system. Examples of this are:

- A system has failed and its data needs to be moved to other systems.
- The work load must be balanced across a number of systems.

In such cases, it is important that the data related to particular application(s) or users be located on an identifiable set of disk drives, so that when these drives are moved, all data of the application(s) or group of users, and no other data, is moved.

Note – The Volume Manager supports a default disk group, `rootdg`, in which all volumes are created if no further specification is given. All commands will default to `rootdg` as well.

To create the disk group `newdg` associated with a disk, do the following:

1. **Type the following to start** `vxdiskadd`**:**

```
vxdiskadd c1t1d0
```

2. **At the following prompt, press Return to continue:**

```
Add or initialize disks
Menu: VolumeManager/Disk/AddDisks

Here is the disk selected.  Output format: [Device_Name]

  c1t1d0

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

3. **At the following prompt, specify the disk group to which the disk should be added (**`newdg`**, in this case):**

```
  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) newdg
```

4. `vxdiskadd` **confirms that no active disk group currently exists with the same name and prompts for confirmation that you really want to create this new disk group**:

```
There is no active disk group named newdg.

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

Type `y` to continue.

5. **At the following prompt, either press Return to accept the default disk name or type a disk name:**

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

6. **When** `vxdiskadd` **asks whether this disk should become a hot-relocation spare, type** `n` **(or press Return):**

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

7. **To continue with the operation, type** `y` **(or press Return) at the following prompt**:

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

  c1t1d0

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

Messages similar to the following should now confirm that this disk is being initialized for VxVM use:

```
  Initializing device c1t1d0.

  Creating a new disk group named newdg containing the disk
  device c1t1d0 with the name newdg01.
```

8. **To see if the disk group was created, type**:

```
vxdisk list
```

```
DEVICE TYPE DISK GROUP STATUS
c0t0d0s2 sliced disk04 rootdg online
c1t0d0s2 sliced disk03 rootdg online
c1t1d0s2 sliced newdg01newdgonline
```

10.4.2 Moving Disk Groups

A disk group can be moved between systems, along with its VxVM objects. In this way, the disk group's configuration is relocated to a new system.

Move a disk group across systems as follows:

1. **Unmount and stop all volumes in the disk group on the first system.**
2. **Deport (disable local access to) the disk group to be moved with the command:**

```
vxdg deport diskgroup
```

3. **Import (enable local access to) the disk group and its disks from the second system with the command:**

```
vxdg import diskgroup
```

4. **After the disk group is imported, start all volumes in the disk group with the command:**

```
vxrecover -g diskgroup -sb
```

10.4.2.1 Renaming Disk Groups

Since only one disk group of a given name can exist per system, you must rename a disk group if you want to move it to a system already containing a disk group with the same name.

Every system running the Volume Manager must have a single `rootdg` disk group. `rootdg` can therefore only be moved across systems if it is renamed.

The following set of steps can be used to temporarily move the `rootdg` disk group from one host to another (for repair work on the root volume, for instance) and then move it back:

1. **On the original host, identify the disk group ID of the root disk group to be imported:**

```
vxdisk -s list
```

This command results in output that includes the following.

```
dgname: rootdg
dgid:   774226267.1025.tweety
```

2. **On the importing host, import and rename the `rootdg` disk group as follows:**

```
vxdg -tC -n newdg_name import diskgroup
```

where `-t` indicates a temporary import name; `-C` clears import locks (and should only be used if you are *certain* that the disks are not being used by any other hosts); `-n` specifies a temporary name for the `rootdg` to be imported (so that it does not conflict with the existing `rootdg`); and *diskgroup* is the ID of the `rootdg` disk group being imported.

If a crash or reboot occurs at this point, the temporarily-imported disk group will become unimported and will require a reimport.

3. **After the necessary work has been done on the imported `rootdg`, deport it back to its original host as follows:**

```
vxdg -h hostname deport diskgroup
```

where *hostname* is the name of the system whose `rootdg` is being returned (the system's name can be confirmed with the command `uname -n`). This command removes the imported `rootdg` from the importing host and returns locks to its original host. The original host will then autoimport its `rootdg` on its next reboot.

10.4.3 Displaying Disk Group Information

To use disk groups, you need to know their names and what disks belong to each group.

Display information on existing disk groups as follows:

```
vxdg list
```

The Volume Manager returns the following listing of current disk groups:

```
NAME    STATE   ID
rootdg  enabled 730344554.1025.tweety
newdg   enabled 731118794.1213.tweety
```

Display more detailed information on a specific disk group (such as `rootdg`) as follows:

```
vxdg list rootdg
```

The Volume Manager returns output such as the following for `rootdg`:

```
Group:     rootdg
dgid:      730344554.1025.tweety
import-id: 0.1
flags:
copies:    nconfig=default nlog=default
config:    seqno=0.94553 permlen=795 free=768 templen=16 loglen=8
config disk c2t2d2s2 copy 1 len=795 disabled
config disk c2t2d3s2 copy 1 len=795 disabled
config disk c2t2d4s2 copy 1 len=795 disabled
config disk c2t3d0s2 copy 1 len=795 state=clean online
log disk c2t2d2s2 copy 1 len=120 disabled
log disk c2t2d3s2 copy 1 len=120 disabled
log disk c2t2d4s2 copy 1 len=120 disabled
log disk c2t3d0s2 copy 1 len=120
```

If you need to verify the disk group ID and name associated with a specific disk (to import the disk group, for example), use the following command:

```
vxdisk -s list devicename
```

This command provides output that includes the following information for the specified disk:

```
dgname: rootdg
dgid:   774226267.1025.tweety
```

CHAPTER **11**

File System CLI Operations

This chapter provides instructions on performing the following file system operations using the command line interface:

Note – Some operations discussed in this chapter can only be performed with the VERITAS File System (VxFS).

11.1 Creating a File System on a Volume

A file system can be created on a volume, but cannot be larger than the volume on which it is created. If you want to use a file system that is expandable and has quick recovery, specify `vxfs` as the file system type.

Create a `vxfs` file system on an existing volume as follows:

```
mkfs [-F vxfs] [generic_options] [-o specific_options] special_file size
```

where `vxfs` is the file system type, *generic_options* are the options common to most file systems, *specific_options* are options specific to the `vxfs` file system, *special_file* is the full path name of the volume on which to create the file system (such as `/dev/vx/rdsk/pubs`), and *size* is the size of the new file system.

For more information about the options and variables available for use with the `mkfs` command, see the `mkfs`(1M) manual page.

To create a file system on a volume called `/dev/vx/rdsk/pubs`, first create a volume called `pubs` using `vxassist`. Refer to Chapter 12, “Volume CLI Operations” for instructions on how to create volumes using `vxassist`.

To create a `vxfs` file system on the newly-created `pubs` volume, type the following:

```
mkfs -F vxfs /dev/vx/rdsk/pubs 12288
```

Note – The size of the file system (in this case 12288 sectors) is the same size or smaller than the volume on which the file system is created.

The following message is displayed:

```
Mkfs: make vxfs file system?
(DEL if wrong)
```

If you have not changed your mind after several seconds, the file system is created, and information similar to the following is displayed:

```
12288 sectors, 6144 blocks of size 1024
1320 inodes, 5294 data blocks, log size 512
1 allocation units of 5630 blocks, 1320 inodes, 5294 data blocks
first allocation unit starts at block 514
overhead per allocation unit is 336 blocks
```

The newly-created file system can now be mounted.

11.2 Mounting a File System

After creating the file system, mount it as follows:

```
mount [-F vxfs][generic_options][-o specific_options] block_special mount_point
```

where *block_special* is a block special device and *mount_point* is the location where the file system will be mounted.

To mount the file system `/dev/vx/dsk/pubs`, type:

```
mount -F vxfs /dev/vx/dsk/pubs /pubs
```

11.3 Unmounting a File System

If you no longer need to access the data in a file system, you can unmount it. Unmount a file system as follows:

```
umount block_special | mount_point
```

where *block_special* is a block special device and *mount_point* is the location where the file system is mounted.

To unmount the file system `/dev/vx/dsk/pubs`, type:

```
umount /dev/vx/dsk/pubs
```

11.4 Resizing a File System

A `vxfs` file system (on a volume) can be resized using the `fsadm` command, as described in this section. Note that the `fsadm` command resizes the file system only.

Alternatively, the `vxresize` command can be used to resize both a file system and its underlying volume to a specified new volume length. Refer to the `vxresize`(1M) manual page for further information.

11.4.1 Extending a File System

If you find that your `vxfs` file system is not large enough, you can increase the length of the file system. If the underlying volume is not already large enough to accommodate the new file system size, you need to increase the size of the volume first (using `vxassist`). The new size of the file system is specified in units of 512-byte blocks (or sectors).

Extend a `vxfs` file system as follows:

```
fsadm [-F vxfs] [-b newsize] [-r rawdev] mount_point
```

where *newsize* is the size to which the file system will grow (in sectors), and *rawdev* can optionally be used to specify the path name of the raw device if it cannot be determined by `fsadm`.

To extend a `vxfs` file system using `/pubs` as a mount point, type:

```
fsadm -F vxfs -b 22528 /pubs
```

11.4.2 Shrinking a File System

If you find that you have allocated more space for your `vxfs` file system than you need, you can decrease the size of the file system.

Note – In cases where data is allocated towards the end of the file system, shrinkage may not be possible and you may have to reorganize the file system. See the *VERITAS File System (VxFS) System Administrator's Guide* for details.

Shrink a `vxfs` file system as follows:

```
fsadm [-F vxfs] [-b newsize] [-r rawdev] mount_point
```

where *newsize* is the size to which the file system will grow (in sectors), and *rawdev* can optionally be used to specify the path name of the raw device if it cannot be determined by `fsadm`.

To shrink a `vxfs` file system to 20480 sectors using `/pubs` as a mount point, type:

```
fsadm -F vxfs -b 20480 /pubs
```

11.5 Displaying Mounted File Systems

It is important to keep track of which file systems are mounted and which are not. This saves users from trying to access unmounted file systems. You can look at the status of mounted file systems by typing:

```
mount
```

When invoked without any options, the mount command displays file system information similar to the following:

```
/ on /dev/root read/write/setuid on Thu May 26 16:58:24 1994
/proc on /proc read/write on Thu May 26 16:58:25 1994
/dev/fd on /dev/fd read/write on Thu May 26 16:58:26 1994
/tmp on /tmp read/write on Thu May 26 16:59:33 1994
/var/tmp on /var/tmp read/write on Thu May 26 16:59:34 1994
```

CHAPTER **12**

Volume CLI Operations

This chapter provides instructions on performing the following volume operations using the command line interface:

Refer to the *Sun StorEdge Volume Manager 2.6 System Administrator's Guide* for additional information and examples on using the command line interface to perform volume operations.

12.1 Creating a Volume

Volumes need to be created to take advantage of the Volume Manager concept of virtual disks. Once a volume exists, a file system can be placed on the volume to organize the disk space with files and directories. Alternatively, applications like databases might be used to organize data on volumes.

The Volume Manager enables you to create volumes with various layout types using the command line interface:

- Concatenated
- Striped
- RAID-5

The `vxassist` command provides the simplest way to create new volumes.

When creating a volume, `vxassist` takes the following form:

```
vxassist make volume_name length [attributes]
```

where `make` is the keyword for volume creation, *volume_name* is a name you give to the volume, and *length* specifies the number of sectors (by default) in the volume. The length can be specified in kilobytes, megabytes, or gigabytes by using a suffix character of `k`, `m`, or `g`, respectively. Refer to the `vxintro`(1M) manual page for more information on specifying units of length when creating volumes. Additional attributes can be specified, as appropriate.

By default, `vxassist` creates volumes in the `rootdg` disk group. Another disk group can be specified by including `-g` *diskgroup* in the `vxassist` command line.

Creating a volume in the disk group `rootdg` creates two device node files that can be used to access the volume:

- `/dev/vx/dsk/`*volume_name* (the block device node for the volume)
- `/dev/vx/rdsk/`*volume_name* (the raw device node for the volume)

For volumes in disk groups other than `rootdg`, these names include the disk group name, as follows:

- `/dev/vx/dsk/`*diskgroup_name/volume_name*
- `/dev/vx/rdsk/`*diskgroup_name/volume_name*

Section 12.1.1, “Creating a Concatenated Volume,” describes the simplest way to create a (default) volume. Later sections describe how to create volumes with specific attributes.

12.1.1 Creating a Concatenated Volume

By default, `vxassist` creates a concatenated volume that uses one or more sections of disk space. On a fragmented disk, this enables you to put together a volume larger than any individual section of free disk space available.

If there is not enough space on a single disk, `vxassist` creates a spanned volume. A spanned volume is a concatenated volume with sections of disk space spread across more than one disk. A spanned volume can be larger than the single largest disk, since it takes space from more than one disk.

12.1.1.1 Creating a Concatenated Volume on Any Disk

If no disk is specified, the Volume Manager selects an appropriate disk on which to create the volume.

Create a concatenated, default volume as follows:

```
vxassist make volume_name length
```

where *volume_name* is a name of the volume and *length* specifies the length of the volume in sectors (unless another unit of size is specified with a suffix character).

To create the volume voldef with a length of 10 megabytes, type:

```
vxassist make voldef 10m
```

12.1.1.2 Creating a Concatenated Volume on a Specific VM Disk

The Volume Manager automatically selects the disk(s) each volume will reside on, unless you specify otherwise. If you want a volume to reside on a specific disk, you must designate that disk for the Volume Manager. More than one disk can be specified.

Create a volume on a specific disk as follows:

```
vxassist make volume_name length diskname [...]
```

To create the volume volspec on disk03, type:

```
vxassist make volspec 3m disk03
```

12.1.2 Creating a Striped Volume

A striped volume contains at least one plex that consists of two or more subdisks located on two or more physical disks. For more information on striping, refer to Chapter 1, “Introduction to the Volume Manager.”

Create a striped volume as follows:

```
vxassist make volume_name length layout=stripe
```

To create the striped volume volzebra, type:

```
vxassist make volzebra 10m layout=stripe
```

This creates a striped volume with the default stripe unit size on the default number of disks.

You can indicate the disks on which the volumes are to be created by specifying the disk names at the end of the command line. For example, to create a 30 megabyte striped volume on three specific disks (disk03, disk04, and disk05), type:

```
vxassist make stripevol 30m layout=stripe disk03 disk04 disk05
```

12.1.3 Creating a RAID-5 Volume

A RAID-5 volume contains a RAID-5 plex that consists of two or more subdisks located on two or more physical disks. Only one RAID-5 plex can exist per volume. A RAID-5 volume may also contain one or more RAID-5 log plexes, which are used to log information about data and parity being written to the volume. For more information on RAID-5 volumes, refer to Chapter 1, “Introduction to the Volume Manager.”

Create a RAID-5 volume as follows:

```
vxassist make volume_name length layout=raid5
```

To create the RAID-5 volume volraid, type:

```
vxassist make volraid 10m layout=raid5
```

This creates a RAID-5 volume with the default stripe unit size on the default number of disks. It also creates a RAID-5 log, by default.

12.2 Removing a Volume

Once a volume is no longer necessary (it is inactive and archived, for example), you can remove the volume and free up the disk space for other uses.

Before removing a volume, make sure you do the following:

1. **Remove all references to the volume.**

2. **If the volume is mounted as a file system, unmount it with the command:**

```
umount /dev/vx/dsk/volume_name
```

3. **If the volume is listed in** `/etc/vfstab`**, remove its entry.**

4. **Make sure that the volume is stopped with the command:**

```
vxvol stop volume_name
```

The `vxvol stop` command stops all Volume Manager activity to the volume.

After following these steps, remove the volume as follows:

```
vxedit -rf rm volume_name
```

To remove the volume `volspan`, type:

```
vxedit -rf rm volspan
```

12.3 Mirroring a Volume

A mirror is a copy of a volume. The mirror copy is not stored on the same disk(s) as the original copy of the volume. Mirroring a volume ensures that the data in that volume will not be lost if one of your disks fails.

12.3.1 Mirroring a New Volume

Create a new volume with a mirror as follows:

```
vxassist make volume_name length layout=mirror
```

To create the mirrored volume, `volmir`, type:

```
vxassist make volmir 5m layout=mirror
```

To create a volume with Dirty Region Logging enabled, create a mirrored volume with a dirty region log as follows:

```
vxassist make volume_name length layout=mirror,log
```

This specifies that the volume layout must be both mirrored and contain a DRL log. `vxassist` creates one log plex per log subdisk, by default.

12.3.2 Mirroring an Existing Volume

Create a mirror for an existing volume as follows:

```
vxassist mirror volume_name
```

The following example creates a mirror of the volume `voltest`:

```
vxassist mirror voltest
```

12.3.3 Mirroring All Volumes

To mirror all existing volumes to available disk space, type:

```
/etc/vx/bin/vxmirror -a
```

You can also configure the Volume Manager to create mirrored volumes by default. To do this, type the command:

```
/etc/vx/bin/vxmirror -d yes
```

If you make this change, you can still make unmirrored volumes by specifying `nmirror=1` as an attribute to the `vxassist` command. For example, to create an unmirrored 20-megabyte volume named `nomirror`, use the command:

```
vxassist make nomirror 20m nmirror=1
```

12.4 Adding a RAID-5 Log

Only one RAID-5 plex can exist per RAID-5 volume. Any additional plexes become RAID-5 log plexes, which are used to log information about data and parity being written to the volume. When a RAID-5 volume is created using `vxassist`, a log plex is created for that volume, by default.

Add a RAID-5 log to an existing volume as follows:

```
vxassist addlog volume_name
```

The following example creates a log for the RAID-5 volume `volraid`:

```
vxassist addlog volraid
```

12.5 Adding a DRL Log

To put Dirty Region Logging into effect for a volume, a log subdisk must be added to that volume and the volume must be mirrored. Only one log subdisk can exist per plex.

Add a DRL log to an existing volume as follows:

```
vxassist addlog volume_name
```

The following example creates a log for the volume `vol03`:

```
vxassist addlog vol03
```

When vxassist is used to add a log subdisk to a volume, a log plex is also created to contain the log subdisk, by default.

Once created, the plex containing a log subdisk can be treated as a regular plex. Data subdisks can be added to the log plex. The log plex and log subdisk can be removed using the same procedures used to remove ordinary plexes and subdisks.

12.6 Removing a Mirror

You may have more than one mirror for each volume. If you find that you have more mirrors than you need, you can remove a plex.

Note – The last valid mirror associated with a volume cannot be removed.

Removing a mirror involves first dissociating the plex from its volume and then removing the plex and any associated subdisks completely.

Dissociate and remove the plex from its volume as follows:

```
vxplex -o rm dis plex_name
```

For example, the following output from vxprint -h shows that volume volmir is mirrored using plexes volmir-01 and volmir-02:

```
Disk group: rootdg

TY  NAME       ASSOC      KSTATE    LENGTH    PLOFFS   STATE    TUTIL0 PUTIL0
v   volmir     fsgen      ENABLED   10240     -        ACTIVE   -      -
pl  volmir-01  volmir     ENABLED   10240     -        ACTIVE   -      -
sd  disk01-03  volmir-01  -         10240     0        -        -      -
pl  volmir-02  volmir     ENABLED   10240     -        ACTIVE   -      -
sd  disk02-02  volmir-02  -         10240     0        -        -      -
```

To disassociate and remove volmir-02 from volmir, type:

```
vxplex -o rm dis volmir-02
```

The output of `vxprint -h` for `volmir` is now:

```
Disk group: rootdg

TY    NAME       ASSOC      KSTATE   LENGTH   PLOFFS   STATE   TUTIL0  PUTIL0
v     volmir     fsgen      ENABLED  10240    -        ACTIVE  -       -
pl    volmir-01  volmir     ENABLED  10240    -        ACTIVE  -       -
sd    disk01-03  volmir-01  -        10240    0        -       -       -
```

12.7 Removing a RAID-5 Log

Removing a RAID-5 log involves first dissociating the log from its volume and then removing the log and any associated subdisks completely.

Dissociate the log from its volume as follows:

```
vxplex -o rm dis plex_name
```

To identify the log plex, use the command:

```
vxprint -ht volume_name
```

where *volume_name* is the name of the RAID-5 volume. This should result in output that lists a plex with the `STATE` field of `LOG`.

To disassociate the log plex `volraid-02` from `volraid`, type:

```
vxplex -o rm dis volraid-02
```

The output of `vxprint -h` for `volraid` is now:

```
Disk group: rootdg

TY   NAME        ASSOC       KSTATE    LENGTH   PLOFFS   STATE   TUTIL0   PUTIL0
v    volraid     raid5       ENABLED   32       -        ACTIVE  -        -
pl   volraid-01  volraid     ENABLED   32       -        ACTIVE  -        -
sd   disk01-03   volraid-01  -         32       0        -       -        -
sd   disk02-02   volraid-01  -         32       0        -       -        -
```

12.8 Extending a Volume

If the volume is not large enough for the amount of data that needs to be stored in it, you need to extend the volume's length.

12.8.1 Extending to a Given Length

Extend a volume *to* a specific length as follows:

```
vxassist growto volume_name length
```

To extend `volcat` to 2000 sectors, type:

```
vxassist growto volcat 2000
```

12.8.2 Extending by a Given Length

Extend a volume *by* a specific length as follows:

```
vxassist growby volume_name length
```

To extend `volcat` by 100 sectors, type:

```
vxassist growby volcat 100
```

12.9 Shrinking a Volume

If you find that your volume is much larger than you really need it to be, you can shrink the volume's size.

Caution – Do not shrink a volume below the size of the file system. If you have a VxFS file system, you can shrink the file system and then shrink the volume. If you do not shrink the file system first, you risk unrecoverable data loss.

12.9.1 Shrinking to a Given Length

Shrink a volume *to* a specific length as follows:

```
vxassist shrinkto volume_name length
```

Make sure you do not shrink the volume below the current size of the file system or database using the volume. This command can be safely used on empty volumes.

To shrink `volcat` to 1300 sectors, type:

```
vxassist shrinkto volcat 1300
```

12.9.2 Shrinking by a Given Length

Shrink a volume *by* a specific length as follows:

```
vxassist shrinkby volume_name length
```

To shrink `volcat` by 300 sectors, type:

```
vxassist shrinkby volcat 300
```

12.10 Preparing a Volume for Online Backup

It is important to make backup copies of your volumes. This provides a copy of the data as it stands at the time of the backup. Backup copies are used to restore volumes lost due to disk failure, or data destroyed due to human error. The Volume Manager enables you to back up volumes with minimal interruption of the volume's availability for users.

Note – The volume backup described here does not apply to RAID-5 volumes.

Backing up a volume with `vxassist` involves the following procedure:

1. **Create a snapshot mirror of the volume to be backed up.**

   The `vxassist snapstart` operation creates a write-only backup mirror, which is attached to and synchronized with the volume to be backed up. When synchronized with the volume, the backup mirror is ready to be used as a snapshot mirror. However, it continues being updated until it is detached during the actual snapshot portion of the procedure. This may take some time, depending on the volume size.

2. **Choose a suitable time to create a snapshot volume.**

   If possible, plan to take the snapshot at a time when users are accessing the volume as little as possible.

3. **Create a snapshot volume that reflects the original volume at the time of the snapshot.**

   The online backup procedure is completed by running the `vxassist snapshot` command on the volume with the snapshot mirror. This operation detaches the finished snapshot mirror, creates a new normal volume, and attaches the snapshot mirror to it. The snapshot then becomes a read-only volume. This step should only take a few minutes.

   The snapshot volume can now be used by backup utilities while the original volume continues to be available for applications and users.

Create a snapshot mirror for a volume as follows:

```
vxassist snapstart volume_name
```

To create a snapshot mirror of a volume called `voldef`, type:

```
vxassist snapstart voldef
```

After creating the snapshot mirror, create a snapshot volume that reflects the original volume as follows:

```
vxassist snapshot volume_name new_volume_name
```

To create a snapshot volume of `voldef`, type:

```
vxassist snapshot voldef snapvol
```

You can now back up the snapshot volume by whatever means you prefer. To avoid wasting space, you can then remove the snapshot volume, which occupies as much space as the original volume.

12.11 Displaying Volume Configuration Information

The `vxprint` command can be used to display information about how a volume is configured.

Display the volume, mirror, and subdisk record information for all volumes as follows:

```
vxprint -ht
```

This displays output such as the following:

```
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     759096729.1025.tweety

dm  disk10     c1t0d0s2 sliced   559      1044400   -
dm  disk20     c2t0d0s2 sliced   559      1044400   -

v   pubs       fsgen    ENABLED  ACTIVE   2288     SELECT     -
pl  pubs-01    pubs     ENABLED  ACTIVE   2288     CONCAT     -        RW
sd  disk10-01  pubs-01  disk10   0        2288     0          c0t0d0   ENA

v   voldef     fsgen    ENABLED  ACTIVE   20480    SELECT     -
pl  voldef-01  voldef   ENABLED  ACTIVE   20480    CONCAT     -        RW
sd  disk10-02  voldef-0 disk10   2288     20480    0          c0t1d0   ENA
```

where `dg` is a disk group, `dm` is a disk, `v` is a volume, `pl` is a plex, and `sd` is a subdisk. The top few lines indicate the headers that match each type of output line that follows. Each volume is listed along with its associated plex(es) and subdisk(s).

Display volume-related information for a specific volume as follows:

```
vxprint -t volume_name
```

To display information about `voldef`, type:

```
vxprint -t voldef
```

This displays output similar to the following:

```
Disk group: rootdg

V   NAMEUSETYPEKSTATESTATELENGTHREADPOLPREFPLEX

v   voldeffsgenENABLEDACTIVE20480SELECT-
```

CHAPTER **13**

Menu Interface Operations

This chapter describes the menu-driven interface known as Volume Manager Support Operations (or `vxdiskadm`). `vxdiskadm` is used to perform physical and logical device administration.

13.1 Starting `vxdiskadm`

To start `vxdiskadm`, type:

```
vxdiskadm
```

`vxdiskadm` starts up and displays its main menu containing the operations available to you.

13.1.1 The vxdiskadm Main Menu

The following is the vxdiskadm main menu:

```
Volume 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
list List disk information

?    Display help about menu
??   Display help about the menuing system
q    Exit from menus

Select an operation to perform
```

- ? can be typed at any time to provide help in using the menu. The output of ? is a list of operations and a definition of each.
- ?? lists inputs that can be used at any prompt.
- q returns you to the main menu if you need to restart a process; however, using q at the main menu level exits the Volume Manager Support Operations.

The vxdiskadm menu provides access to the following operations:

- Add or initialize one or more disks

  This menu operation enables you to add formatted disks to the system. SCSI disks are already formatted. For other disks, see the manufacturer's documentation for formatting instructions. You are prompted for the disk device(s). You can specify the disk group to which the disk(s) should be added; if none is selected, the disk is held as a spare to be used for future operations or disk replacements without needing to be initialized at that time. You can also specify that selected disks be marked as hot-relocation spares for a disk group. If the disk has not been initialized already, the disk is partitioned and initialized for use with the Volume Manager.

- Encapsulate one or more disks

  This menu operation is used to bring a disk that was added to the system before installing the Volume Manager under Volume Manager control. You are prompted for the disk device(s), disk group, and disk name(s). The disk is added to the specified disk group. The disk is examined to search for partitions that are used for file systems or other purposes. Volumes are created to replace disk partitions as the means of accessing the existing data. If the encapsulation cannot determine the purpose of a partition automatically, you are asked what to do with it. You can choose to replace the partition with a volume, leave the partition alone, or remove the partition.

  You must reboot the system if any partitions are being used for mounted file systems or for running applications. You may have to modify application configuration files to use volumes, rather than direct disk devices, to access the disk partitions. File system mount information is adjusted automatically.

- Remove a disk

  This menu operation is used to remove a disk from a disk group. You are prompted for the name of a disk to remove. You cannot remove a disk if any volumes use storage on that disk. If any volumes are using storage on the disk, you have the option of asking the Volume Manager to move that storage to other disks in the disk group.

Note – You cannot remove the last disk in a disk group using this operation. If you want to use all the remaining disks in a disk group for some purpose, you should disable (deport) the disk group. You will then be able to reuse the disks. `rootdg` cannot, however, be deported.

- Remove a disk for replacement

  This menu operation is used to remove a physical disk from a disk group, while retaining the disk name. This changes the state for the named disk to `removed`. If there are any initialized disks that are not part of a disk group, you are given the option of using one of these disks as a replacement.

- Replace a failed or removed disk

  This menu operation is used to specify a replacement disk for a disk that you removed with the `Remove a disk for replacement` menu entry, or one that failed during use. You are 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.

- Mirror volumes on a disk

This menu operation is used to mirror volumes on a disk. These volumes can be mirrored to another disk with available space. Creating mirror copies of volumes in this way protects against data loss in case of disk failure. Volumes that are already mirrored or that are comprised of more than one subdisk will not be mirrored with this operation. Mirroring volumes from the boot disk will produce a disk that can be used as an alternate boot disk.

- Move volumes from a disk

  This menu operation is used to move any volumes (or parts of a volume) that are using a disk onto other disks. Use this menu operation immediately prior to removing a disk, either permanently or for replacement.

Note – Simply moving volumes off a disk, without also removing the disk, does not prevent other volumes from being moved onto the disk by future operations.

- Enable access to (import) a disk group

  This menu operation enables access by this system to 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. Then, move the disks from the deported disk group to the other system and enable (import) the disk group there. You are prompted for the disk group name.

- Disable access to (deport) a disk group

  This menu operation disables 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 are prompted for the name of a disk group. You are 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.

- Enable (online) a disk device

  If you move a disk from one system to another during normal system operation, the Volume Manager will not recognize the disk automatically. Use this menu operation to tell the Volume Manager to scan the disk to identify it, and to determine if this disk is part of a disk group. Also, use this operation to re-enable access to a disk that was disabled by either the disk group deport operation or the disk device disable (offline) operation.

- Disable (offline) a disk device

  This menu operation is used to disable all access to a disk device through the Volume 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 that some systems do not support disks that can be removed from a system during normal operation. On such systems, the offline operation is seldom useful.

- Mark a disk as a spare for a disk group

  This menu operation is used to reserve a disk as an automatic replacement disk (for hot-relocation) in case another disk in the disk group should fail.

- Turn off the spare flag on a disk

  This menu operation is used to free hot-relocation spare disks for use as regular VM disks.

- List disk information

  This menu operation is used to display a list of disks attached to your system. This also lists removed or failed disks.

  You can also use this operation to list detailed information for a particular disk. This information includes the disk group of which the disk is a member, even if that disk group is currently disabled.

13.2 Disk Operations

This section describes the disk operations available with `vxdiskadm`.

When performing disk administration, it is important that you recognize the difference between a *device name* and a *disk name.*

The *device name* (sometimes referred to as *devname* or *disk access name*) is the location of the disk. The syntax of a device name is `c#t#d#s#`, where:

- `c#` is the number of the controller to which the disk drive is attached.
- `t#` is the number of the target disk on that controller.
- `d#` is the number of the disk (or UNIX partition).
- `s#` is the number of the disk slice.

Note – While working with the `vxdiskadm` interface, the slice portion of the device name (`s#`) should not be entered. Only `c#t#d#` should be used.

The full path name of a device is `/dev/dsk/`*devicename*. In this document, only the device name is listed and `/dev/dsk` is assumed. An example of a device name is `c0t0d0s2`.

The *disk name* (sometimes referred to as *disk media name*) is an administrative name for the disk, such as `disk01`. If you do not assign a disk name, the disk name defaults to `disk##` if the disk is being added to `rootdg` (where ## is a sequence number). Otherwise, the default disk name is *groupname##*, where *groupname* is the name of the disk group to which the disk is added.

13.2.1 Bringing Physical Disks Under Volume Manager Control

When you add a disk to a system that is running Volume Manager, you need to put the disk under Volume Manager control so that the Volume Manager can control the space allocation on the disk.

Unless another disk group is specified, VxVM places new disks in the default disk group, `rootdg`. Instructions on creating additional disk groups are provided later in this chapter.

The method by which you place a disk under Volume Manager control depends on the circumstances:

- If the disk is new, it will need to be initialized and placed under Volume Manager control (see Section 13.2.1.1, "Placing a Disk Under VxVM Control").
- If the disk is not needed immediately, it can be initialized (but not added to a disk group) and reserved for future use (see Section 13.2.2, "Adding a Disk for Future Use").
- If the disk was previously initialized for future Volume Manager use, it can be reinitialized and placed under VxVM control (see Section 13.2.3, "Reinitializing a Disk for VxVM Use").
- If the disk was previously in use, but not under Volume Manager control, you may want to preserve existing data on the disk while still letting Volume Manager take control of the disk. This can be accomplished using encapsulation (see Section 13.2.1.3, "Encapsulating a Disk for VxVM Use").
- Multiple disks on one or more controllers can be placed under VxVM control simultaneously. Depending on the circumstances, some of the disks may be handled differently to others during this operation. See Section 13.2.1.2, "Placing Multiple Disks Under VxVM Control."

When initializing or encapsulating multiple disks at once, it is possible to exclude certain disks or certain controllers. To exclude disks, list the names of the disks to be excluded in the file `/etc/vx/disks.exclude` before the initialization or encapsulation. Similarly, you can exclude all disks on specific controllers from initialization or encapsulation by listing those controllers in the file `/etc/vx/cntrls.exclude`.

The sections that follow provide detailed examples of how to use `vxdiskadm` to place disks under VxVM control in various ways and circumstances.

Note – A disk must be formatted (using the `format` command, for example) before it can be placed under Volume Manager control.

If you attempt to place an unformatted disk under Volume Manager control through `vxdiskadm`, the initialization will begin as normal, but will quit with a message that informs you that the disk does not appear to be valid and may not be formatted. If this happens, you need to format the disk properly and then attempt to place the disk under VxVM control again.

13.2.1.1 Placing a Disk Under VxVM Control

This section describes how to place a formatted disk under Volume Manager control. This disk might be brand new or might have been used outside VxVM before.

Note – Initialization does not preserve any data on disks.

Initialize a single disk for Volume Manager use as follows:

1. **Select menu item** `1` **(**`Add or initialize one or more disks`**) from the** `vxdiskadm` **main menu.**

2. **At the following prompt, type the disk device name of the disk to be added to VxVM control (or type** `list` **for a list of disks):**

```
Add or initialize disks
Menu: VolumeManager/Disk/AddDisks

Use this operation to add one or more disks to a disk group. You can add the
selected disks to an existing disk group or to a new disk group that will be
created as a part of the operation. The selected disks may also be added to a
disk group as spares. The selected disks may also be initialized without adding
them to a disk group leaving the disks available for use as replacement disks.

More than one disk or pattern may be entered at the prompt. Here are some disk
selection examples:

  all:         all disks
  c3 c4t2:     all disks on both controller 3 and controller 4, target 2
  c3t4d0:      a single disk

Select disk devices to add:
[<pattern-list>,all,list,q,?] list
```

<*pattern-list*> can be a single disk, or a series of disks and/or controllers (with optional targets). If <*pattern-list*> consists of multiple items, those items must be separated by white space.

If you type `list` at the prompt, `vxdiskadm` displays a list of the disks available to the system, followed by a prompt at which you should type the device name of the disk to be added:

```
DEVICE       DISK         GROUP        STATUS
c0t0d0       disk01       rootdg       online
c0t1d0       disk02       rootdg       online
c0t2d0       disk03       rootdg       online
c0t3d0       -            -            online
c1t0d0       disk10       rootdg       online
c1t0d1       -            -            error
.
.
.
c3t0d0       -            -            error
c3t1d0       disk33       rootdg       online
c3t2d0       disk34       rootdg       online
c3t3d0       disk35       rootdg       online

Select disk devices to add:
[<pattern-list>,all,list,q,?] c1t0d1
```

All disks attached to the system are recognized by the Volume Manager and displayed here.

The word `error` in the `STATUS` line tells you that a disk has not yet been added to Volume Manager control. These disks may or may not have been initialized before. The disks that are listed with a disk name and disk group cannot be used for this operation, as they are already under Volume Manager control.

3. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
  Here is the disk selected.  Output format: [Device_Name]

  c1t2d0

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

4. **At the following prompt, specify the disk group to which the disk should be added or press Return to accept** `rootdg`**:**

```
  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)
```

5. **At the following prompt, either press Return to accept the default disk name or type a disk name:**

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

6. **When** `vxdiskadm` **asks whether this disk should become a hot-relocation spare, type** `n` **(or press Return):**

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

7. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
  The selected disks will be added to the disk group rootdg with
  default disk names.

  c1t2d0

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

8. **If you are sure that there is no data on this disk, type** `n` **to avoid encapsulation. When** `vxdiskadm` **asks if you want to initialize the disk instead, type** `y`**:**

```
  The following disk device has a valid VTOC, but does not appear to have
  been initialized for the Volume 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.
  Output format: [Device_Name]

  c1t2d0

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

  c1t2d0

Instead of encapsulating, initialize? [y,n,q,?] (default: n) y
```

Messages similar to the following should now confirm that disk `c1t2d0` is being placed under Volume Manager control.

```
  Initializing device c1t2d0.

  Adding disk device c1t2d0 to disk group rootdg with disk
  name disk39.
```

9. **At the following prompt, indicate whether you want to continue to initialize more disks (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Add or initialize other disks? [y,n,q,?] (default: n)
```

13.2.1.2 Placing Multiple Disks Under VxVM Control

This section describes how to place multiple disks under Volume Manager control simultaneously. The set of disks might consist of all disks on the system, all disks on a controller, selected disks, or a combination thereof.

Depending on the circumstances, some of the specified disks might be handled differently to others. For example, some may be initialized, while others may be encapsulated.

Note – Initialization does not preserve any data on disks.

When initializing or encapsulating multiple disks at once, it is possible to exclude certain disks or certain controllers. To exclude disks, list the names of the disks to be excluded in the file `/etc/vx/disks.exclude` before the initialization or encapsulation. Similarly, you can exclude all disks on specific controllers from initialization or encapsulation by listing those controllers in the file `/etc/vx/cntrls.exclude`.

Place multiple disks under VxVM control at once as follows:

1. **Select menu item** `1` **(**`Add or initialize one or more disks`**) from the** `vxdiskadm` **main menu.**
2. **At the following prompt, type the *pattern-list* for the disks to be added to VxVM control. In this case, type** `c3` **to indicate all disks on controller 3:**

```
Add or initialize disks
Menu: VolumeManager/Disk/AddDisks

Use this operation to add one or more disks to a disk group. You can add the
selected disks to an existing disk group or to a new disk group that will be
created as a part of the operation. The selected disks may also be added to a
disk group as spares. The selected disks may also be initialized without adding
them to a disk group leaving the disks available for use as replacement disks.

More than one disk or pattern may be entered at the prompt. Here are some disk
selection examples:

  all:        all disks
  c3 c4t2:    all disks on both controller 3 and controller 4, target 2
  c3t4d0:     a single disk

Select disk devices to add:
[<pattern-list>,all,list,q,?] c3
```

<*pattern-list*> can be a single disk, or a series of disks and/or controllers (with optional targets). If <*pattern-list*> consists of multiple items, those items must be separated by white space.

If you do not know the address (device name) of the disk to be added, type `list` at the prompt for a complete listing of available disks.

3. **To continue the operation, type** `y` **(or press Return) at the following prompt:**

```
  Here are the disks selected.  Output format: [Device_Name]

  c3t0d0 c3t1d0 c3t2d0 c3t3d0

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

4. **To add these disks to the default disk group,** `rootdg`**, type** `y` **(or press Return) at the following prompt:**

```
  You can choose to add these disks to an existing disk group, a
  new disk group, or you can leave these disks 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 disks available for future use, specify a disk group
  name of “none”.

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

5. **To enable** `vxdiskadm` **to use default disk names for each of the disks, type** `y` **(or Press Return) at the following prompt:**

```
Use default disk names for these disks? [y,n,q,?] (default: y) y
```

6. **At the following prompt, type** `n` **to indicate that these disks should not be used as hot-relocation spares:**

```
Add disks as spare disks for rootdg? [y,n,q,?] (default: n) n
```

7. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
  The selected disks will be added to the disk group rootdg with
  default disk names.

  c3t0d0 c3t1d0 c3t2d0 c3t3d0

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

8. **The following prompt lists any disks that have already been initialized for VxVM use; type** `Y` **to indicate that all of these disks should be used now:**

```
  The following disk devices appear to have been initialized already.
  The disks are currently available as replacement disks.
  Output format: [Device_Name]

  c3t1d0 c3t2d0 c3t3d0

Use these devices? [Y,N,S(elect),q,?] (default: Y) Y
```

Note that this prompt enables you to indicate "yes" or "no" for *all* of these disks (`Y` or `N`) or to select how to handle each of these disks on an individual basis (`S`).

If you are sure that you want to reinitialize all of these disks, type `Y` at the following prompt:

```
  The following disks you selected for use appear to already have
  been initialized for the Volume Manager.  If you are certain the
  disks already have been initialized for the Volume Manager, then
  you do not need to reinitialize these disk devices.
  Output format: [Device_Name]

  c3t1d0 c3t2d0 c3t3d0

Reinitialize these devices? [Y,N,S(elect),q,?] (default: Y) Y
```

9. `vxdiskadm` **now indicates that one of the disks on controller 3 is a candidate for encapsulation; type** `y` **(or press Return) to encapsulate this disk:**

```
  The following disk device has a valid VTOC, but does not appear to have
  been initialized for the Volume 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.
  Output format: [Device_Name]

  c3t0d0

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

Encapsulation enables you to add an active disk to Volume Manager control and preserve the data on that disk.

Note – Disk encapsulation requires that you reboot the system and may require a few subsequent reboots. You will be prompted for these operations, as necessary.

`vxdiskadm` now confirms those disks that are being initialized and added to VxVM control with messages similar to the following:

```
Initializing device c3t1d0.

Initializing device c3t2d0.

Initializing device c3t3d0.

Adding disk device c3t1d0 to disk group rootdg with disk
name disk33.

Adding disk device c3t2d0 to disk group rootdg with disk
name disk34.

Adding disk device c3t3d0 to disk group rootdg with disk
name disk35.
```

10. `vxdiskadm` **then confirms any disks that have been selected for encapsulation and prompts you for permission to proceed with the encapsulation; type** `y` **(or press Return) to continue encapsulation:**

```
  The following disk has been selected for encapsulation.
  Output format: [Device_Name]

  c3t0d0

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

`vxdiskadm` now displays an encapsulation status and informs you that you must perform a shutdown and reboot as soon as possible:

```
 The disk device c3t0d0 will be encapsulated and added to the disk group
  rootdg with the disk name disk38.

  The first stage of encapsulation has completed successfully.  You
  should now reboot your system at the earliest possible opportunity.

  The encapsulation will require two or three reboots which will happen
  automatically after the next reboot.  To reboot execute the command:

shutdown -g0 -y -i6

  This will update the /etc/vfstab file so that volume devices are
  used to mount the file systems on this disk device.  You will need
  to update any other references such as backup scripts, databases,
  or manually created swap devices.
```

11. **At the following prompt, indicate whether you want to continue to initialize more disks (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Add or initialize other disks? [y,n,q,?] (default: n)
```

13.2.1.3 Encapsulating a Disk for VxVM Use

This section describes how to encapsulate a disk for VxVM use. Encapsulation preserves any existing data on the disk when the disk is placed under VxVM control.

To reduce the chance of encapsulation failure, make sure that the disk:

- Has a small amount of free space (at the beginning or end of the disk) that does not belong to any partition
- Has two free partitions
- Has an `s2` slice that represents the whole disk

When encapsulating the boot (root) disk, the swap partition should be tagged as `swap` so that it will be possible to dump to that partition later.

Before encapsulating your boot disk, the EEPROM variable `use-nvramrc?` must be set to `true` if you want to take advantage of the Volume Manager boot disk aliases to identify the mirror of the boot disk if a replacement is needed. If this variable is set to `false`, you will need to determine which disks are bootable yourself. You can set this variable to `true` as follows:

```
eeprom use-nvramrc?=true
```

Encapsulate a disk for VxVM use as follows:

1. **Select menu item** `2` **(**`Encapsulate one or more disks`**) from the** `vxdiskadm` **main menu.**
2. **At the following prompt, type the disk device name for the disks to be encapsulated:**

```
Encapsulate one or more disks
Menu: VolumeManager/Disk/Encapsulate

  Use this operation to convert one or more disks to use the Volume  Manager.
  This adds the disks to a disk group and replaces existing partitions
  with volumes.  Disk encapsulation requires a reboot for the changes
  to take effect.

  More than one disk or pattern may be entered at the prompt.  Here are
  some disk selection examples:

  all:         all disks
  c3 c4t2:     all disks on both controller 3 and controller 4, target 2
  c3t4d0:      a single disk

Select disk devices to encapsulate:
[<pattern-list>,all,list,q,?] c2t5d0
```

<*pattern-list*> can be a single disk, or a series of disks and/or controllers (with optional targets). If <*pattern-list*> consists of multiple items, those items must be separated by white space.

If you do not know the address (device name) of the disk to be encapsulated, type `l` or `list` at the prompt for a complete listing of available disks.

3. **To continue the operation, type `y` (or press Return) at the following prompt:**

```
  Here is the disk selected.  Output format: [Device_Name]

  c2t5d0

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

4. **To add the disk to the default disk group, `rootdg`, press Return at the following prompt:**

```
  You can choose to add this disk to an existing disk group or to
  a new disk group.  To create a new disk group, select a disk group
  name that does not yet exist.

Which disk group [<group>,list,q,?] (default: rootdg)
```

5. **At the following prompt, either press Return to accept the default disk name or type a disk name:**

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

6. **To continue with the operation, type `y` (or press Return) at the following prompt:**

```
  The selected disks will be encapsulated and added to the rootdg
  disk group with default disk names.

  c2t5d0

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

7. **To confirm that encapsulation should proceed, type `y` (or press Return) at the following prompt:**

```
  The following disk has been selected for encapsulation.  Output format:
[Device_Name]

  c2t5d0

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

A message similar to the following confirms that the disk is being encapsulated for VxVM use and tells you that a reboot is needed:

```
 The disk device c2t5d0 will be encapsulated and added to the disk
  group  rootdg with the disk name disk01.

  The c2t5d0 disk has been configured for encapsulation.

  The first stage of encapsulation has completed successfully.  You
  should now reboot your system at the earliest possible opportunity.

  The encapsulation will require two or three reboots which will happen
  automatically after the next reboot.  To reboot execute the command:

shutdown -g0 -y -i6

  This will update the /etc/vfstab file so that volume devices are
  used to mount the file systems on this disk device.  You will need
  to update any other references such as backup scripts, databases,
  or manually created swap devices.
```

8. **At the following prompt, indicate whether you want to encapsulate more disks (`y`) or return to the `vxdiskadm` main menu (`n`):**

```
Encapsulate other disks? [y,n,q,?] (default: n) n
```

Under some circumstances, encapsulation of a disk may fail. Encapsulation often fails because there is not enough free space available on the disk to accommodate VxVM. If this happens, the procedure outlined above will end abruptly with an error message such as the following:

```
  The disk device c2t5d0 will be encapsulated and added to the disk group
  rootdg with the disk name disk01.

  The encapsulation operation failed with the following error:

  It is not possible to encapsulate c2t5d0, for the following reason:
     <vxvm:vxslicer: ERROR: Unsupported disk layout.>

Hit RETURN to continue.
```

13.2.2 Adding a Disk for Future Use

If you want to add a disk to Volume Manager control for future use, follow the steps outlined in Section 13.2.1, “Bringing Physical Disks Under Volume Manager Control.” However, when you are asked to name a disk group, type `none` instead of selecting `rootdg` or typing in a disk group name. The disk is then initialized as before, but is reserved for use at some later time. It cannot be used until it is added to a disk group.

Note that this type of “spare disk” should not be confused with a hot-relocation spare disk.

13.2.3 Reinitializing a Disk for VxVM Use

This section describes how to reinitialize a disk that has already been initialized for VxVM use at some time in the past.

If the disk you want to add has been used before, but not with VxVM, you can do *either* of the following:

- Encapsulate the disk and preserve its information (see Section 13.2.1.3, “Encapsulating a Disk for VxVM Use”).
- Reinitialize the disk, allowing the Volume Manager to configure the disk for VxVM. Note that reinitialization does not preserve data on the disk. If you want to have the disk reinitialized, make sure that the disk does not contain data that should be preserved.

Reinitialize a disk for VxVM use as follows:

1. **Select menu item** `1` **(**`Add or initialize one or more disks`**) from the** `vxdiskadm` **main menu.**

2. **At the following prompt, type the disk device name of the disk to be added to VxVM control:**

```
Add or initialize disks
Menu: VolumeManager/Disk/AddDisks

Use this operation to add one or more disks to a disk group. You can add the
selected disks to an existing disk group or to a new disk group that will be
created as a part of the operation. The selected disks may also be added to a
disk group as spares. The selected disks may also be initialized without adding
them to a disk group leaving the disks available for use as replacement disks.

More than one disk or pattern may be entered at the prompt. Here are some disk
selection examples:

  all:          all disks
  c3 c4t2:      all disks on both controller 3 and controller 4, target 2
  c3t4d0:       a single disk

Select disk devices to add:
[<pattern-list>,all,list,q,?] c1t3d0
```

<*pattern-list*> can be a single disk, or a series of disks and/or controllers (with optional targets). If <*pattern-list*> consists of multiple items, those items must be separated by white space.

If you do not know the address (device name) of the disk to be added, type `l` or `list` at the prompt for a complete listing of available disks.

3. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
  Here is the disk selected.  Output format: [Device_Name]

  c1t3d0

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

4. **At the following prompt, specify the disk group to which the disk should be added or press Return to accept** `rootdg`**:**

```
  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)
```

5. **At the following prompt, either press Return to accept the default disk name or type a disk name:**

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

6. **When** `vxdiskadm` **asks whether this disk should become a hot-relocation spare, type** `n` **(or press Return):**

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

7. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
  The selected disks will be added to the disk group rootdg with
  default disk names.

  c1t3d0

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

8. **The following prompt indicates that this disk has been previously initialized for future VxVM use; type** `y` **to confirm that you now want to use this disk:**

```
  The following disk device appears to have been initialized already.
  The disk is currently available as a replacement disk.
  Output format: [Device_Name]

  c1t3d0

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

9. **To reinitialize the disk, type** `y` **(or press Return) at the following prompt:**

```
  The following disk you selected for use appears to already have
  been initialized for the Volume Manager.  If you are certain the
  disk has already been initialized for the Volume Manager, then you
  do not need to reinitialize the disk device.
  Output format: [Device_Name]

  c1t3d0

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

Messages similar to the following should now confirm that this disk is being reinitialized for VxVM use:

```
Initializing device c1t3d0.

Adding disk device c1t3d0 to disk group rootdg with disk
name disk40.
```

10. **At the following prompt, indicate whether you want to continue to initialize more disks (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Add or initialize other disks? [y,n,q,?] (default: n)
```

13.2.4 Adding a VM Disk to the Hot-Relocation Pool

Hot-relocation is the ability of a system to automatically react to I/O failure by relocating redundant subdisks to other disks and restoring the affected VxVM objects and data. If a disk has already been designated as a spare in the disk group, the subdisks from the failed disk are relocated to the spare disk. Otherwise, any suitable free space in the disk group is used. Refer to Chapter 1, “Introduction to the Volume Manager” for more information.

To designate a disk as a spare, do the following:

1. **Select menu item** `12` **(**`Mark a disk as a spare for a disk group`**) from the** `vxdiskadm` **main menu.**

2. **At the following prompt, type a disk name (such as** `disk01`**):**

```
Mark a disk as a spare for a disk group
Menu: VolumeManager/Disk/MarkSpareDisk

  Use this operation to mark a disk as a spare for 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,?] disk01
```

`vxdiskadm` displays the following confirmation:

```
  Marking of disk01 in rootdg as a spare disk is complete.
```

3. **At the following prompt, indicate whether you want to add more disks as spares (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Mark another disk as a spare? [y,n,q,?] (default: n)
```

Any VM disk in this disk group can now use this disk as a spare in the event of a failure. If a disk fails, hot-relocation should automatically occur (if possible). You should be notified of the failure and relocation via electronic mail. After successful relocation, you may want to replace the failed disk.

13.2.5 Removing a VM Disk From the Hot-Relocation Pool

While a disk is designated as a spare, the space on that disk is not used as free space for the creation of VxVM objects within its disk group. If necessary, you can free a spare disk for general use by removing it from the pool of hot-relocation disks.

To verify which disks are currently designated as spares, select the `list` menu item from the `vxdiskadm` main menu. Disks that are spares should be listed with the `spare` flag.

To remove a disk from the hot-relocation pool, do the following:

1. **Select menu item** `13` **(**`Turn off the spare flag on a disk`**) from the** `vxdiskadm` **main menu.**

2. **At the following prompt, type the name of a spare disk (such as `disk01`):**

```
Turn off the spare flag on a disk
Menu: VolumeManager/Disk/UnmarkSpareDisk

  Use this operation to turn off the spare flag on a disk.
  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,?] disk01
```

`vxdiskadm` displays the following confirmation:

```
  Disk disk01 in rootdg no longer marked as a spare disk.
```

3. **At the following prompt, indicate whether you want to disable more spare disks (`y`) or return to the `vxdiskadm` main menu (`n`):**

```
Turn-off spare flag on another disk? [y,n,q,?] (default: n)
```

13.2.6 Displaying Information on Physical Disks

Displaying disk information shows you which disks are initialized, to which disk groups they belong, and the disk status. The `list` command displays device names for all recognized disks, the disk names, the disk group names associated with each disk, and the status of each disk.

To display disk information, do the following:

1. **Select `list` from the `vxdiskadm` main menu.**

2. **At the following screen, type the address of the disk you want to see or type** `all` **for a list of all disks:**

```
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)
```

- If you type `all`, the Volume Manager displays the following information:

```
DEVICE    DISK       GROUP     STATUS
c0t0d0    c0t0d0s2   rootdg    online
c1t0d0    disk01     rootdg    online
c1t1d0    -          -         online

Device to list in detail [<address>,none,q,?] (default: none)
```

- If you type the address of the device for which you want information, complete disk information (including the device name, the type of disk, and information about the public and private partitions of the disk) is shown.

  Once you have examined this information, press Return to return to the main menu.

13.2.7 Removing a Physical Disk

Before removing a disk, make sure it contains no data, any data is no longer needed, or the data can be moved to other disks. Then remove the disk as follows:

1. **Select menu item** `3` **(**`Remove a disk`**) from the** `vxdiskadm` **main menu.**

Note – You must disable the disk group before you can remove the last disk in that group.

---

2. **At the following prompt, type the disk name of the disk to be removed:**

```
Remove a disk
Menu: VolumeManager/Disk/RemoveDisk

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,?] disk01
```

3. **If there are any volumes on the disk, the Volume Manager asks you whether they should be evacuated from the disk. If you want to keep the volumes, answer** `y`**. Otherwise, answer** `n`**.**

4. **At the following verification prompt, press Return to continue:**

```
Requested operation is to remove disk disk01 from group rootdg.

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

`vxdiskadm` removes the disk from the disk group and displays the following success message:

```
Removal of disk disk01 is complete.
```

You can now remove the disk or leave it on your system as a replacement.

5. **At the following prompt, indicate whether you want to remove other disks (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Remove another disk? [y,n,q,?] (default: n)
```

13.2.8 Disabling a Physical Disk (Taking a Physical Disk Offline)

Occasionally, you may need to take a disk offline. If the disk is corrupted, you need to take it offline and remove it. You may be moving the physical disk device to another location to be connected to another system. To take a disk offline, first remove it from its disk group, then do the following:

1. **Select menu item** `11` **(**`Disable (offline) a disk device`**) from the** `vxdiskadm` **main menu.**
2. **At the following prompt, type the address of the disk you want to disable:**

```
Disable (offline) a disk device
Menu: VolumeManager/Disk/OfflineDisk

Use this menu operation to disable all access to a disk device
by the Volume 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,?] c1t1d0
```

`vxdiskadm` disables the specified disk.

3. **At the following prompt, indicate whether you want to disable another device (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Disable another device? [y,n,q,?] (default: n)
```

13.2.9 Enabling a Physical Disk

If you move a disk from one system to another during normal system operation, the Volume Manager does not recognize the disk automatically. Use the enable disk operation to tell the Volume Manager to scan the disk to identify the disk, and to determine if this disk is part of a disk group. Also, use this operation to re-enable access to a disk that was disabled by either the disk group deport operation or the disk device disable (offline) operation. To enable a disk, do the following:

1. **Select menu item** `10` **(**`Enable (online) a disk device`**) from the** `vxdiskadm` **main menu.**

2. **At the following prompt, type the device name of the disk to be enabled (or type** `list` **for a list of devices):**

```
Enable (online) a disk device
Menu: VolumeManager/Disk/OnlineDisk

Use this operation to enable access to a disk that was disabled
with the "Disable (offline) a disk device" operation.

You can also use this operation to re-scan a disk that may have
been changed outside of the Volume Manager. For example, if a disk
is shared between two systems, the Volume Manager running on the
other system may have changed the disk. If so, you can use this
operation to re-scan the disk.

NOTE: Many vxdiskadm operations re-scan disks without user
      intervention. This will eliminate most needs to online a
      disk directly, except when the disk is directly offlined.

Select a disk device to enable [<address>,list,q,?] c1t1d0
```

`vxdiskadm` enables the specified device.

3. **At the following prompt, indicate whether you want to enable another device (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Enable another device? [y,n,q,?] (default: n)
```

13.2.10 Replacing a Physical Disk

If a disk fails, you need to replace that disk with another. This operation requires disabling and removing the failed disk and installing a new disk in its place. To replace a disk, do the following:

1. **Select menu item** `4` **(**`Remove a disk for replacement`**) from the** `vxdiskadm` **main menu.**

2. **At the following prompt, type the name of the disk to be replaced (or type** `list` **for a list of disks):**

```
Remove a disk for replacement
Menu: VolumeManager/Disk/RemoveForReplace

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,?] disk02
```

Additional screens display if there are volumes associated with the disk you want to remove. You must decide whether to keep the data associated with the volumes or to allow that data to be lost when the disk is replaced. Answer any prompts accordingly.

3. **At the following prompt, either select the device name of the replacement disk (from the list provided) or press Return to choose the default disk:**

```
The following devices are available as replacements:
     c1t1d0

You can choose one of these disks now, to replace disk02.
Select "none" if you do not wish to select a replacement disk.

Choose a device, or select "none"
[<device>,none,q,?] (default: c1t1d0)
```

4. **At the following prompt, press Return to continue:**

```
Requested operation is to remove disk disk02 from group rootdg.
The removed disk will be replaced with disk device c1t1d0.

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

`vxdiskadm` displays the following success screens:

```
Removal of disk disk02 completed successfully.

Proceeding to replace disk02 with device c1t1d0.
```

```
Disk replacement completed successfully.
```

5. **At the following prompt, indicate whether you want to remove another disk (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Remove another disk? [y,n,q,?] (default: n)
```

13.2.11 Replacing a Failed or Removed Disk

Disks can be removed and then replaced later. To do this, use menu item 4 (`Remove a disk for replacement`) to remove a disk, then do the following:

1. **Select menu item** `5` **(**`Replace a failed or removed disk`**) from the** `vxdiskadm` **main menu.**
2. **Select the disk name of the disk to be replaced:**

```
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,?] disk02
```

3. `vxdiskadm` **displays the device names of the disk devices available for use as replacement disks; type the device name of the device of your choice or press Return to select the default device:**
4. **At the following prompt, press Return to replace the disk:**

   `vxdiskadm` displays the following success screen:

```
The following devices are available as replacements:
     c1t0d0s2 c1t1d0s2

You can choose one of these disks to replace disk02.
Choose "none" to initialize another disk to replace disk02.

Choose a device, or select "none"
[<device>,none,q,?] (default: c1t0d0s2)
```

```
The requested operation is to use the initialized device c1t0d0s2
to replace the removed or failed disk disk02 in disk group rootdg.

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

```
Replacement of disk disk02 in group rootdg with disk device
c1t0d0s2 completed successfully.
```

5. **At the following prompt, indicate whether you want to replace another disk (`y`) or return to the `vxdiskadm` main menu (`n`):**

```
Replace another disk? [y,n,q,?] (default: n)
```

13.3 Volume Operations

This section describes the volume operations that can be performed with the Volume Manager Support Operations. These volume operations are used to protect data in case of device failures.

13.3.1 Mirroring Volumes on a VM Disk

Mirroring the volumes on a VM disk gives you one or more copies of your volumes in another disk location. By creating mirror copies of your volumes, you protect your system against loss of data in case of a disk failure. You can use this operation on your root disk to make a second copy of the boot information available on an alternate disk. This will enable you to boot your system even if your root disk is corrupted.

To mirror volumes on a disk, make sure that the target disk has an equal or greater amount of space as the originating disk and then do the following:

1. **Select menu item** `6` **(**`Mirror volumes on a disk`**) from the** `vxdiskadm` **main menu.**
2. **At the following prompt, type the disk name of the disk whose volumes you want to mirror:**

```
Mirror volumes on a disk
Menu: VolumeManager/Disk/Mirror

This operation can be used to mirror volumes on a disk. These
volumes can be mirrored onto another disk or onto any
available disk space. Volumes will not be mirrored if they are
already mirrored. Also, volumes that are comprised of more than
one subdisk will not be mirrored.

Mirroring volumes from the boot disk will produce a disk that can
be used as an alternate boot disk.

Enter disk name [<disk>,list,q,?] disk02
```

3. **At the following prompt, type the target disk name (this disk must be the same size or larger than the originating disk):**

```
You can choose to mirror volumes from disk disk02 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) disk01
```

Note – Be sure to always specify the destination disk when you are creating an alternate root disk. Otherwise, the Volume Manager will select a disk to be the alternate root disk; however, your system may not be able to boot from that disk.

---

4. **At the following prompt, press Return to make the mirror:**

```
The requested operation is to mirror all volumes on disk disk02
in disk group rootdg onto available disk space on disk disk01.

NOTE: This operation can take a long time to complete.

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

`vxdiskadm` displays the status of the mirroring operation:

```
Mirror volume voltest-bk00 ...

Mirroring of disk disk01 is complete.
```

5. **At the following prompt, indicate whether you want to mirror volumes on another disk (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Mirror volumes on another disk? [y,n,q,?] (default: n)
```

13.3.2 Moving Volumes from a VM Disk

Before you disable or remove a disk, you may want to move the data from that disk to other disks on the system. To do this, make sure that the target disks have sufficient space, then do the following:

1. **Select menu item** `7` **(**`Move volumes from a disk`**) from the** `vxdiskadm` **main menu.**

2. **At the following prompt, type the disk name of the disk whose volumes you want to move:**

```
Move volumes from a disk
Menu: VolumeManager/Disk/Evacuate

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,?] disk01
```

After the following screen is displayed, you can optionally specify a list of disks to which the volume(s) should be moved.

```
You can now specify a list of disks to move onto.  Specify a list
of disk media names (e.g., disk01) 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.
```

3. **At the following prompt, press Return to move the volumes:**

```
Requested operation is to move all volumes from disk disk01 in
group rootdg.

NOTE: This operation can take a long time to complete.

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

As the volumes are moved from the disk, `vxdiskadm` displays the status of the operation:

```
Move volume voltest ...
Move volume voltest-bk00 ...
```

When the volumes have all been moved, `vxdiskadm` displays the following success screen:

```
Evacuation of disk disk01 is complete.
```

4. **At the following prompt, indicate whether you want to move volumes from another disk (`y`) or return to the `vxdiskadm` main menu (`n`):**

```
Move volumes from another disk? [y,n,q,?] (default: n)
```

13.4 Disk Group Operations

The following sections describe disk group operations available with the Volume Manager Support operations.

13.4.1 Adding a Disk to a Disk Group

You may want to add a new disk to an already established disk group. Perhaps the current disks have insufficient space for the project or work group requirements, especially if these requirements have changed. You can add a disk to a disk group by following the steps required to add a disk. Refer to Section 13.2.1, "Bringing Physical Disks Under Volume Manager Control."

13.4.2 Creating a Disk Group

There can be situations in which all data related to a particular set of applications of a particular group of users needs to be made accessible on another system. Examples of this are:

- A system has failed and its data must be moved to other systems.
- The work load must be balanced across a number of systems.

In such cases, it is important that the data related to particular application(s) or users be located on an identifiable set of disk drives. When these disks are moved, all data belonging to the application(s) or users on these disks, and no other data, is moved.

Note – The Volume Manager supports a default disk group, `rootdg`, in which all volumes are created if no further specification is given. All commands will default to `rootdg` as well.

A disk group can only be created along with a disk. A disk group must have at least one disk associated with it.

If you need to create a disk group in addition to `rootdg`, do the following:

1. **Select menu item** `1` **(**`Add or initialize one or more disks`**) from the** `vxdiskadm` **main menu.**
2. **At the following prompt, type the disk device name of the disk to be added to VxVM control:**

```
Add or initialize disks
Menu: VolumeManager/Disk/AddDisks

Use this operation to add one or more disks to a disk group. You can add the
selected disks to an existing disk group or to a new disk group that will be
created as a part of the operation. The selected disks may also be added to a
disk group as spares. The selected disks may also be initialized without adding
them to a disk group leaving the disks available for use as replacement disks.

More than one disk or pattern may be entered at the prompt. Here are some disk
selection examples:

  all:         all disks
  c3 c4t2:     all disks on both controller 3 and controller 4, target 2
  c3t4d0:      a single disk

Select disk devices to add:
[<pattern-list>,all,list,q,?] c1t2d0
```

<*pattern-list*> can be a single disk, or a series of disks and/or controllers (with optional targets). If <*pattern-list*> consists of multiple items, those items must be separated by white space.

If you do not know the address (device name) of the disk to be added, type `l` or `list` at the prompt for a listing of all disks.

3. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

```
  Here is the disk selected.  Output format: [Device_Name]

  c1t2d0

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

4. **At the following prompt, specify the disk group to which the disk should be added (**`anotherdg`**, in this case):**

```
  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) anotherdg
```

5. `vxdiskadm` **confirms that no active disk group currently exists with the same name and prompts for confirmation that you really want to create this new disk group:**

```
There is no active disk group named anotherdg.

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

6. **At the following prompt, either press Return to accept the default disk name or type a disk name:**

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

7. **When** `vxdiskadm` **asks whether this disk should become a hot-relocation spare, type** `n` **(or press Return):**

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

8. **To continue with the operation, type** `y` **(or press Return) at the following prompt:**

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

  c1t2d0

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

9. **The following prompt is displayed if this disk was previously initialized for future VxVM use; type** `y` **to confirm that you now want to use this disk:**

```
  The following disk device appears to have been initialized already.
  The disk is currently available as a replacement disk.
  Output format: [Device_Name]

  c1t2d0

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

If you are sure you want to reinitialize the disk, type `y` (or press Return) at the following prompt:

```
  The following disk you selected for use appears to already have
  been initialized for the Volume Manager.  If you are certain the
  disk has already been initialized for the Volume Manager, then you
  do not need to reinitialize the disk device.
  Output format: [Device_Name]

  c1t2d0

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

Messages similar to the following should now confirm that this disk is being reinitialized for VxVM use:

```
  Initializing device c1t2d0.

  Creating a new disk group named anotherdg containing the disk
  device c1t2d0 with the name another01.
```

10. **At the following prompt, indicate whether you want to continue to initialize more disks (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Add or initialize other disks? [y,n,q,?] (default: n)
```

13.4.3 Deporting a Disk Group

Use the deport disk group operation to disable 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.

To deport a disk group, do the following:

1. **Select menu item** `9` **(**`Remove access to (deport) a disk group`**) from the** `vxdiskadm` **main menu.**
2. **At the following prompt, type the name of the disk group to be deported:**

```
Remove access to (deport) a disk group
Menu: VolumeManager/Disk/DeportDiskGroup

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) newdg
```

3. **At the following prompt, type** `y` **if you intend to remove the disks in this disk group:**

```
The requested operation is to disable access to the removable
disk group named newdg. This disk group is stored on the
following disks:
     newdg01 on device c1t1d0s2

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)
```

4. **At the following prompt, press Return to continue with the operation:**

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

Once the disk group is deported, `vxdiskadm` displays the following screen:

```
Removal of disk group newdg was successful.
```

5. **At the following prompt, indicate whether you want to disable another disk group (**`y`**) or return to the** `vxdiskadm` **main menu (**`n`**):**

```
Disable another disk group? [y,n,q,?] (default: n)
```

13.4.4 Importing a Disk Group

Use this menu operation to enable access by this system to 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, then move the disk between systems and enable (import) the disk group.

To import a disk group, do the following:

1. **Select menu item** `8` **(**`Enable access to (import) a disk group`**) from the** `vxdiskadm` **main menu.**

2. **At the following prompt, type the name of the disk group to import:**

```
Enable access to (import) a disk group
Menu: VolumeManager/Disk/EnableDiskGroup

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)
newdg
```

Once the import is complete, `vxdiskadm` displays the following success screen:

```
The import of newdg was successful.
```

3. **At the following prompt, indicate whether you want to import another disk group (`y`) or return to the `vxdiskadm` main menu (`n`):**

```
Select another disk group? [y,n,q,?] (default: n)
```

13.5 Exiting the Volume Manager Support Operations

When you have completed all of your disk administration activities, exit the Volume Manager Support Operations.

To exit from the Volume Manager Support Operations, select menu item `q` from the `vxdiskadm` main menu.

APPENDIX **A**

X Resources for the Visual Administrator

This appendix lists X resources that can be used to configure the Visual Administrator according to personal preferences and system requirements. The default values specified here correspond to the defaults that have been compiled into the Visual Administrator. Preferences specified in the system's `app-defaults` file may change these defaults.

The files in `/opt/SUNWvxva/app-defaults` contain a subset of resources that can be used to customize some aspects of the Visual Administrator, such as fonts and colors.

The X resources and associated preferences can be specified in the user's `.Xdefaults` file. Refer to the X Window System documentation on X resources for further information.

Visual Administrator entries in the `.Xdefaults` file should take the following form:

`vxva*`*resource*: *value*

For example, the color used to represent a disabled icon can be altered from the default color (grey) to orange by editing the `.Xdefaults` file to include the following line:

```
vxva*disabledPixel: orange
```

Resources for the Visual Administrator can also be specified for a single session only by invoking the Visual Administrator as follows:

```
vxva -xrm "vxva*resource:     value"
```

The following are the resources and their default values, which can be changed according to user preferences. The resources are listed to the left with their default values to the right. Each resource-value pair is prefaced by a brief description.

A.1 Color Resources

The following resources apply only when the Visual Administrator is run on a color monitor. The current default settings are listed below.

The basic background colors will affect the entire background, unless you specify an override value. The default background and foreground colors are as follows:

background	white
foreground	black

The colors used for all menus and menu bars, except the command window:

menubar*background	SteelBlue
menubar*foreground	white

The background color used under the main window buttons (both values must be set to the same color):

swindow*XmDrawingArea.backgrou nd	#33648B
swindow*row_column.background	SteelBlue4

The background color used for the main window menu bar:

RootWin.menubar*background	SteelBlue
RootWin.menubar*foreground	White

The colors used for the command window menu bar:

cmd_main.cmd_menubar*background	SteelBlue
cmd_main.cmd_menubar*foreground	White

The basic outline colors used for the command window:

cmd_main*background	#4F94CD
cmd_main*cmd_pane*foreground	White

The colors used for the output text area in the command window:

cmd_main*cmd_pane*XmText.backgr ound	#63B8FF
cmd_main*cmd_pane*XmText.foregr ound	Black

The colors used for the history and the running command lists in the command window:

cmd_main*cmd_pane*XmScrolledWin dow.XmList.background	#63B8FF
cmd_main*cmd_pane*XmScrolledWin dow.XmList.foreground	Black

The color used for the menu bar in the view windows:

TopLevelShell.XmForm.menubar*fo reground	White
TopLevelShell.XmForm.menubar*ba ckground	#63B8FF

The color used for the background in the view windows and scroll bars:

TopLevelShell.XmForm.ViewScroll.XmDrawingArea.backgro und	#d9f0ff
TopLevelShell.XmForm.ViewScroll.XmScrollBar.backgroun d	White
TopLevelShell.XmForm.ViewScroll.background	White

The background color used for forms:

TopLevelShell.propertiesForm*background	#4F94CD

The color used for the foreground for text forms:

TopLevelShell.propertiesForm*foreground	White

The highlight color used behind form text fields and toggle buttons. For consistency, use the same color:

TopLevelShell.propertiesForm.XmForm*XmText.background	#7EC0EE
TopLevelShell.propertiesForm.XmForm.Toggle.XmToggleBu tton*background	#7EC0EE

The highlight color used behind the file system information property sheet text fields and toggle buttons. For consistency, use the same color:

TopLevelShell.propertiesForm.XmForm.FsMenuForm*XmText .background	#7EC0EE
TopLevelShell.propertiesForm.XmForm.XmForm.Toggle.XmT oggleButton.background	#7EC0EE

The default color used for a volume in a view:

TopLevelShell.XmForm.ViewScroll.XmDrawingArea*VolWidg etClass.background	#99cdeb

The default color used for a VM disk in a view:

TopLevelShell.XmForm.ViewScrol l.XmDrawingArea*VmdiskWidgetCl ass.background	#99cdeb

The default color used for a physical disk in a view:

TopLevelShell.XmForm.ViewScroll.XmDrawingArea*DiskWidgetClass.background	#99cdeb

The default color used for a plex in a view:

TopLevelShell.XmForm.ViewScroll.XmDrawingArea*PlexWidgetClass.background	#a3dbfa

The default color used for a subdisk in a view:

TopLevelShell.XmForm.ViewScroll.XmDrawingArea*SubdiskWidgetClass.background	#c2efff

The default color used for a physical disk partition which is not associated with a disk group:

TopLevelShell.XmForm.ViewScroll.XmDrawingArea*PartWidgetClass.background	#c2efff

The colors used for the notice popup windows:

Notice_popup*background	SteelBlue
Notice_popup*foreground	White

The color used for icons that have been selected:

selectedPixel	green

The color used for icons that are disabled and cannot be used by the Visual Administrator. For example, the color to be used for detached plexes:

```
disabledPixel                    gray
```

The color used for icons that have been selected when a Visual Administrator error occurs (incorrectly selected icons, for example):

```
alarmPixel                        red
```

The color used for subdisk icons that are free (unassociated) when “Show Free Subdisks” has been turned on:

```
freesdPixel                      yellow
```

The color used for icons that are projecting (displaying object relationships) when “Icon Projection” has been turned on for that icon or a related icon:

```
projectPixel                     yellow
```

The color used for icons that have a *low* usage level (as defined in the Analysis Properties Form [TABLE 3-6]) when analysis has been turned on for that icon or a related icon:

```
lowPixel                         green
```

The color used for icons that have a *medium* usage level (as defined in the Analysis Properties Form [TABLE 3-6]) when analysis has been turned on for that icon or a related icon:

```
midPixel                         yellow
```

The color used for icons that have a *high* usage level (as defined in the Analysis Properties Form [TABLE 3-6]) when analysis has been turned on for that icon or a related icon:

```
highPixel                        red
```

The color used to indicate whether the Visual Administrator should run in monochrome mode. When `True` is specified, the Visual Administrator is forced to operate in monochrome (black and white) mode, whether or not a color monitor is being used:

mono	False

A.2 Monochrome Resources

The following resources only apply when the Visual Administrator is run on a monochrome monitor.

The color in which all foreground items are displayed. This typically applies to icon outlines and text:

foreground	black

The color that serves as the background for all windows in the Visual Administrator:

background	white

The bitmap pattern for icons that have been selected:

selectedPixmap	gray3

The bitmap pattern for icons that are disabled and cannot be used by the Visual Administrator (detached plexes, for example):

disabledPixmap	stripe4

The bitmap pattern for icons that have been selected when a Visual Administrator error occurs (incorrectly selected icons, for example):

alarmPixmap	gray1

The bitmap pattern for subdisk icons that are free (unassociated) when “Show Free Subdisks” has been turned on:

```
freesdPixmap                      root_weave
```

The bitmap pattern for icons that are projecting (displaying object relationships) when “Icon Projection” has been turned on for that icon or a related icon:

```
projectPixmap                     root_weave
```

The bitmap pattern for icons that have a *low* usage level (as defined in the Analysis Properties Form) when analysis has been turned on for that icon or a related icon:

```
lowPixmap                         cross_weave
```

The bitmap pattern for icons that have a *medium* usage level (as defined in the Analysis Properties Form) when analysis has been turned on for that icon or a related icon:

```
midPixmap                         root_weave
```

The bitmap pattern for icons that have a *high* usage level (as defined in the Analysis Properties Form) when analysis has been turned on for that icon or a related icon:

```
highPixmap                        wide_weave
```

A.3 Icon Resources

The following resources relate to the Visual Administrator icons.

When `True` is specified, volume icons will be minimized when created, by default:

```
volumeMinimizeIcons               False
```

When `True` is specified, plex icons will be minimized when created, by default. This feature is useful to display structures within volumes, but to hide details about the subdisk structure that makes up the plex:

```
plexMinimizeIcons                False
```

When `True` is specified, disk icons will be minimized when created, by default:

```
diskMinimizeIcons                False
```

When `True` is specified, physical disk icons will be minimized when created, by default:

```
phyDiskMinimizeIcons             False
```

When `True` is specified, icons selected for an operation are automatically de-selected when the operation completes. If set to `False`, icons remain selected until the user decides to deselect them, making it possible to perform multiple operations on the same set of selected icons:

```
autoDeselect                     True
```

A.4 Miscellaneous Resources

The following are miscellaneous resources.

When `True` is specified, the Visual Administrator will be run in tutorial mode whether or not the `-t` command line option has been specified. This function is useful for users who are learning to use the Visual Administrator:

```
tutorial                         False
```

The title of the Visual Administrator's main window:

```
title                            "Visual Administrator"
```

Specify the number of command silos supported. A *command silo* is a set of sequentially dependent commands (like file system create, followed by file system mount). A larger number of silos supports a larger number of concurrent operations that can be run, but also requires the Visual Administrator to use more memory:

```
commandSilos                      50
```

Specify the number of commands that the Visual Administrator should remember and display in the history portion of the Command Info Window:

```
commandHistorySize                20
```

Specify the name of the disk group to be popped up by default when the Visual Administrator is run:

```
defaultViewWindow                 NONE
```

Specify how often, in seconds, the Visual Administrator should check the system mount table to accurately display information about mounted file systems:

```
chkMntptInterval                  5
```

When `True` is specified, the Visual Administrator will force the mouse to be remapped so that it can be used with a two button mouse:

```
twoButtonMouse                    False
```

Changes the height of the Visual Administrator forms and the menu bars. For small resolution displays, this value can be set to `0`:

```
marginHeight                      1
```

Changes the width of the Visual Administrator forms and menu bars. For small resolution displays, this value can be set to `0`:

```
marginWidth                       1
```

The following example shows you how to change the number of columns from its default value of `6` to `12` using a resource:

```
vxva*viewname*numColumns          12
```

Substitute *viewname* with World, `rootdg`, volumes, or user created view names.

A.5 Font Settings

This section lists the existing default font settings.

The following affects the font used in views of plexes, volumes, and subdisks:

```
vxva.fontList                     -*-lucidatypewriter-medium-r-
                                  *-*-*-120-*-*-*-*-*-*
```

The default font used for all menu bars:

```
menubar*fontList                  -*-helvetica-bold-r-normal-*-
                                  *-140-*-*-*-*-*-*
```

The font used for the command window menu bar:

```
cmd_main.cmd_menubar*fontList     -*-helvetica-bold-r-normal-*-
                                  *-140-*-*-*-*-*-*
```

The font used for the labels on the main window buttons:

```
swindow*row_column*fontList        -*-helvetica-medium-r-*-*-*-
                                  140-*-*-*-*-*-*
```

The font used for the form buttons:

```
TopLevelShell.XmForm.XmFrame*f    -*-helvetica-bold-r-normal-*-
ontList                           *-140-*-*-*-*-*-*
```

Glossary

analysis The method used by the Visual Administrator to monitor Volume Manager disk, subdisk, and volume activity.

array A collection of disks logically arranged into an object. Arrays tend to provide benefits such as redundancy or improved performance.

associate The process by which a relationship is established between Volume Manager objects. For example, a subdisk that has been created and defined as having a starting point within a plex is referred to as being *associated* with that plex.

associated plex A plex associated with a volume.

associated subdisk A subdisk associated with a plex.

atomic operation An operation that either succeeds completely or fails and leaves everything as it was before the operation was started. If the operation succeeds, all aspects of the operation take effect at once and the intermediate states of change are invisible. If any aspect of the operation fails, then the operation aborts without leaving partial changes.

attached A state in which a VM object is both associated with another object and enabled for use.

block The minimum unit of data transfer to a disk or array.

boot disk A disk used for booting purposes. This disk may be under VxVM control.

button Either a button on a mouse or a ‘‘button’’ on a window or form.

cascading menu A sub-menu accessed through menu options from the pull-down menus. An arrow at the end of a menu selection indicates the availability of a cascading menu.

check box The icon used to indicate or change the setting of optional controls (for example, default settings). The check box is shaded or contains a check mark to indicate its selection.

column A set of one or more subdisks within a striped plex. Striping is achieved by allocating data alternately and evenly across the columns within a plex.

concatenation A layout style characterized by subdisks that are sequentially and contiguously arranged within the plex address space.

configuration database A set of records containing detailed information on existing Volume Manager objects (such as disk and volume attributes). A single copy of a configuration database is called a configuration copy.

cylinder The iconic representation of disks or volumes within the Visual Administrator.

data stripe This represents the usable data portion of a stripe and is equal to the stripe minus the parity region.

detached A state in which a VM object is associated with another object, but not enabled for use.

detached plex A plex that is inaccessible for reads and writes, but is still associated with the volume. The last plex in a volume is not allowed to be detached.

device name The device name or address used to access a physical disk, such as `c0t0d0s2`. The `c#t#d#s#` syntax identifies the controller, target address, disk, and partition.

Dirty Region Logging The procedure by which the Volume Manager monitors and logs modifications to a plex. A bitmap of changed regions is kept in an associated subdisk called a *log subdisk*.

disk A collection of read / write data blocks that are indexed and can be quickly and randomly accessed. Each disk on a system is given a unique ID that can be used to identify the disk, even if it is moved.

disk access name The device name or address used to access a physical disk, such as `c0t0d0s2`. The `c#t#d#s#` syntax identifies the controller, target address, disk, and partition. The term *device name* can also be used to refer to the disk access name.

disk access record A configuration record used to specify the access path to a particular disk. Each disk access record contains a name and type, and may include some type-specific information.

disk group A collection of disks that share a common configuration. A disk group configuration is a set of records containing detailed information on existing Volume Manager objects (such as disk and volume attributes) and their relationships. Each disk group has an administrator-assigned name and an internally defined unique ID. The root disk group (`rootdg`) is a special private disk group that always exists.

disk media name An administrative name chosen for the disk, such as `disk03`. The term *disk name c*an also be used to refer to the disk media name.

disk media record A configuration record that identifies a particular disk by an administrative name and a unique ID.

dissociate The process by which any link that exists between two Volume Manager objects is removed. For example, dissociating a subdisk from a plex removes the subdisk from the plex and adds the subdisk to the free space pool.

dissociated plex A plex dissociated from a volume.

dissociated subdisk A subdisk dissociated from a plex.

drag and drop The icon manipulation method whereby the icon is selected using the LEFT mouse button, moved to another location in the configuration, and dropped there by releasing the mouse button.

encapsulation A way of placing a disk under Volume Manager control that preserves any existing data on the disk. With encapsulation, existing partitions are converted to volumes and existing file systems are mounted on volumes.

file system A collection of files organized together into a structure. The UNIX file system is a hierarchical structure consisting of directories and files.

free space An area of a disk under Volume Manager control that is not currently allocated to any volume or reserved for use by any Volume Manager object.

free subdisk A subdisk that is not associated with any plex and has an empty `putil[0]` field.

hot-relocation A technique of automatically restoring redundancy and access to mirrored and RAID-5 volumes when a disk fails. This is done by relocating the affected subdisks to disks designated as spares and/or free space in the same disk group.

icon The graphical representation of the Visual Administrator system configuration entities.

iconify The act of turning a window into an icon, or changing the shape and view of a Visual Administrator object icon.

log plex A plex used to store a RAID-5 log.

log subdisk A subdisk used to log recent disk activity using a dirty region log. In the Visual Administrator, log subdisks are represented iconically as rectangles with double borders.

menu bar The bar across the top of each window that contains the menu options for that window.

mirror A duplicate copy of a volume and the data therein. There can be several mirrors per volume. In the Visual Administrator, mirrors are represented iconically as rectangles with heavy borders. The terms *mirror* and *plex* are used synonymously.

mirroring A layout technique that mirrors the contents of a volume onto multiple plexes. Each plex duplicates the data stored on the volume, but the plexes themselves may have different layouts.

object An entity that is defined to and recognized internally by the Volume Manager. The Volume Manager objects are: volume, plex, subdisk, disk, and disk group.

parity A calculated value that can be used to reconstruct data after a failure. While data is being written to a RAID-5 volume, parity is also calculated by performing an *exclusive OR* (XOR) procedure on data. The resulting parity is then written to the volume. If a portion of a RAID-5 volume fails, the data that was on that portion of the failed volume can be recreated from the remaining data and the parity.

parity stripe unit A RAID-5 volume storage region that contains parity information. The data contained in the parity stripe unit can be used to help reconstruct regions of a RAID-5 volume that are missing because of I/O or disk failures.

partition The standard division of a disk device. In the Visual Administrator, partitions are represented iconically as rectangles within physical disk icons. The terms *partition* and *slice* are sometimes used synonymously.

physical disk The underlying storage device, which may or may not be under Volume Manager control. In the Visual Administrator, physical disks are represented iconically as cylinders labeled PD.

plex A duplicate copy of a volume and the data therein. There can be several plexes per volume. Each plex is one copy of the volume with which it is associated. In the Visual Administrator, plexes are represented iconically as rectangles with heavy borders. The terms *plex* and *mirror* are used synonymously.

pop up To open a window.

private region A region of a physical disk used to store private, structured Volume Manager information. The *private region* contains a disk header, a table of contents, and a configuration database. The table of contents maps the contents of the disk. The disk header contains a disk ID. All data in the private region is duplicated for extra reliability.

public region A region of a physical disk managed by the Volume Manager that contains available space and is used for allocating subdisks.

projection The Visual Administrator method used to illustrate the relationships between icons representing Volume Manager objects. When a particular icon is selected for projection, all icons related to it are highlighted (either with color or a bitmap pattern).

pull-down menu The menu selections accessed through choices located in the menu bar of each window.

RAID A Redundant Array of Inexpensive Disks (RAID) is a disk array set up with part of the combined storage capacity used for storing duplicate information about the data stored in that array. This makes it possible to regenerate the data if a disk failure occurs.

radio buttons A set of buttons, only one of which can be selected at any given time. These buttons toggle on or off.

root disk The disk containing the root file system. This disk may be under VxVM control.

root disk group A special private disk group that always exists on the system. The root disk group is named `rootdg`.

root file system The initial file system mounted as part of the UNIX kernel startup sequence.

root partition The disk region on which the root file system resides.

root volume The VxVM volume that contains the root file system, if such a volume is designated by the system configuration.

rootability The ability to place the `root` file system and the `swap` device under Volume Manager control. The resulting volumes can then be mirrored to provide redundancy and allow recovery in the event of disk failure.

sector A unit of size, which can vary between systems. A sector is commonly 512 bytes.

select-operate The icon manipulation method that generally begins by selecting an icon using the LEFT or MIDDLE mouse button. The LEFT button is then used to access the available menu options, and an operation is applied to the selected object(s).

slice The standard division of a logical disk device. The terms *partition* and *slice* are sometimes used synonymously.

spanning A layout technique that permits a file system or database too large to fit on a single disk to span across multiple physical disks.

sparse mirror A plex that is not as long as the volume or that has “holes” (regions of the plex that do not have a backing subdisk).

stripe unit Equally-sized areas that are allocated alternately on the subdisks (within columns) of each striped plex. In an array, this is a set of logically contiguous blocks that exist on each disk before allocations are made from the next disk in the array.

stripe unit size The size of each stripe unit. The default stripe unit size is 64 kilobytes (or 16 kilobytes for RAID-5 stripe units). A *stripe unit size* may also be referred to as a *stripe width.*

stripe A set of stripe units that occupy the same positions across a series of columns.

striping A layout technique that spreads data across several physical disks using stripes. The data is allocated alternately to the stripe columns within a plex.

subdisk A consecutive set of contiguous disk blocks that form a logical disk segment. Subdisks are associated with plexes to form volumes. In the Visual Administrator, subdisks are represented iconically as small rectangles within VM disk or plex icons.

swap area A disk region used to hold copies of memory pages swapped out by the system pager process.

swap volume A VxVM volume that is configured for use as a swap area.

transaction A set of configuration changes that succeed or fail as a group, rather than individually. Transactions are used internally to maintain consistent configurations.

VM disk A disk that is both under Volume Manager control and assigned to a disk group. VM disks are sometimes referred to as *Volume Manager disks* or simply *disks*. In the Visual Administrator, VM disks are represented iconically as cylinders labeled D.

view window A special Visual Administrator window that display icons representing all or a subset of the objects currently known to the Visual Administrator. View windows permit the manipulation of physical and logical views of the mass storage subsystem. A set of default views always exists; users have the option of creating additional user-created views.

views subwindow A smaller window within the Visual Administrator root window that contains view buttons representing existing views.

Visual Administrator main window The Visual Administrator main window through which the user accesses views and menu options. The Visual Administrator main window is the first interactive window that the user sees when the Visual Administrator is started. The Visual Administrator main window is also known as the Visual Administrator *root window*.

volume A virtual disk, which represents an addressable range of disk blocks used by applications such as file systems or databases. A volume can be composed of one or more plexes. In the Visual Administrator, volumes are represented iconically as cylinders labeled V.

window menu button The box generally located in the upper left-hand corner of a window that controls the physical properties of that window.

Index

D

E

F

H

W

X