


VERITAS Volume Manager™ 4.0

User's Guide - VERITAS Enterprise Administrator

Solaris

N10011F



Disclaimer

The information contained in this publication is subject to change without notice. VERITAS Software Corporation makes no warranty of any kind with regard to this manual, including, but not limited to, the implied warranties of merchantability and fitness for a particular purpose. VERITAS Software Corporation shall not be liable for errors contained herein or for incidental or consequential damages in connection with the furnishing, performance, or use of this manual.

Copyright

Copyright © 2001-2004 VERITAS Software Corporation. All rights reserved. VERITAS, the VERITAS logo, and all other VERITAS product names and slogans are trademarks or registered trademarks of VERITAS Software Corporation. VERITAS and the VERITAS Logo Reg. U.S. Pat & Tm. Off. Other product names and/or slogans mentioned herein may be trademarks or registered trademarks of their respective companies.

Data Encryption Standard (DES) Copyright

Copyright © 1990 Dennis Ferguson. All rights reserved.

Commercial use is permitted only if products that are derived from or include this software are made available for purchase and/or use in Canada. Otherwise, redistribution and use in source and binary forms are permitted.

Copyright 1985, 1986, 1987, 1988, 1990 by the Massachusetts Institute of Technology. All rights reserved.

Export of this software from the United States of America may require a specific license from the United States Government. It is the responsibility of any person or organization contemplating export to obtain such a license before exporting.

WITHIN THAT CONSTRAINT, permission to use, copy, modify, and distribute this software and its documentation for any purpose and without fee is hereby granted, provided that the above copyright notice appear in all copies and that both that copyright notice and this permission notice appear in supporting documentation, and that the name of M.I.T. not be used in advertising or publicity pertaining to distribution of the software without specific, written prior permission. M.I.T. makes no representations about the suitability of this software for any purpose. It is provided as is without express or implied warranty.

VERITAS Software Corporation
350 Ellis Street
Mountain View, CA 94043
USA
Phone 650-527-8000
Fax 650-527-2908
<http://www.veritas.com>



Contents

Preface	xix
Audience	xix
Organization	xix
Related Documents	xx
Getting Help	xxi
Downloading and Running VRTSexplorer	xxii
Unique Message Number	xxii
Chapter 1. Introducing the VEA GUI	1
VxVM Features	1
Main Window	2
The Object Tree and Grid	3
The Menu Bar	4
Customizing Table Headers	5
The Toolbar	5
The Status Area	6
Task Roadmap	6
Setup Tasks	6
Maintenance Tasks	7
Repair/Recovery Tasks	8
Chapter 2. Getting Started with VxVM VEA	11
Setting Up Your System	12
Starting VEA	12
Viewing Objects and Object Properties	12
Tree and Grid View	13



Disk Usage Legend	13
Displaying Disk Information	13
Disk View Window	14
Moving Disks in the Disk View	17
Mapping Disks to Volumes	17
Displaying Volume Information	17
Volume View Window	18
Volume Layout Details Window	19
Volume to Disk Mapping Window	19
Navigating the Volume to Disk Mapping Window	21
Mapping Volumes to Disks	22
Dissociated Plex View Window	22
Object Properties Window	22
Navigating the Object Properties Window	23
Multiple Selections	24
Accessing Tasks	24
Viewing Task History	24
Task Throttling	25
Clearing Finished Tasks	25
Viewing Task Properties	26
Command Log	26
Using the Search Feature	26
Set Up a Search Query	26
Use a Saved Query to Search and Display Information	30
View the Results of a Previously Saved Query	31
Edit and Remove Search Queries	31
Performance Tuning	32
Real-time Statistics Collection and Display	32
Select Data Gathering Options	32
Viewing Real-Time Statistics	34

Load Balancing	35
Disk Failure or Removal	35
Disk Phaseout	36
Setting Preferences	36
Show Dialog Top Information Panel	36
Launch Task Dialogs as Modal	36
Show Rules Page	37
Show Display Name of Capability	37
Chapter 3. Disk Tasks	39
Managing Disks	39
Managing Disk Groups	40
Disk Task Roadmap	41
Adding a Disk to VxVM	42
Adding a Disk to a Disk Group	42
Scanning Disks	43
Reserving a Disk	44
Renaming a Disk	45
Onlining a Disk	46
Mirroring a Disk	47
Evacuating a Disk	48
Disconnecting a Disk	49
Replacing a Disk	50
Recovering Volumes on a Disk	51
Removing a Disk from a Disk Group	52
Forcing Initialization of a Disk	53
Resizing Disks	54
Setting Disk Usage	55
Viewing the Table of Disks	56
Viewing Disk Properties	57

Disabling a Controller	58
Enabling a Controller	58
Renaming an Enclosure	59
Creating a Disk Group	60
Upgrading a Disk Group Version	62
Renaming Disk Groups	63
Deporting Disk Groups	64
Importing Disk Groups	65
Recovering Volumes in Disk Groups	66
Commands for Hot Relocation	67
About Hot Spare and Hot Relocation	67
Undo Hot Relocation	68
Clear Hot Relocation Info	69
Customize Disks for Hot Relocation	69
Destroying a Disk Group	71
Moving Objects Between Disk Groups	72
Splitting A Disk Group	74
Joining Disk Groups	76
Viewing the Table of Disk Groups	77
Viewing Disk Group Properties	78
Organizing Disk Groups	78
Chapter 4. Volume Tasks	79
Volume Task Roadmap	80
VxVM Volumes	82
Creating Volumes	86
Creating a VxVM Volume	87
Selecting Your Own Disks	87
VxVM Selecting the Disks	89
Creating a VxVM ISP Volume	91

Create a VxVM ISP Volume Using a New Template	91
Create a VxVM ISP Volume from an Existing User Template	92
Concatenated	94
Striped	95
RAID-5	96
Mirrored	97
Resizing a Volume	98
Resizing and Deleting Multiple VxVM ISP Volumes	99
Renaming a Volume	99
Transforming an ISP Volume	101
Adding and Removing Columns in a VxVM ISP Volume	102
Changing the VxVM ISP Column Count	102
Setting Volume Usage	103
Changing a VxVM Volume Layout	103
Adding a Mirror to a Volume	106
Adding a Log to a Volume	108
Stopping and Recovering a Volume	110
Recovering a Volume	110
Preparing to Restore a Volume From Backup	111
Removing a Mirror From a Volume	112
Removing a Log From a Volume	114
Deleting a Volume	115
Attaching Mirrors	116
Detaching a Mirror	116
Moving, Splitting, and Joining Subdisks	117
Moving a Subdisk	117
Sub-disk Properties	119
Splitting a Subdisk	119
Joining a Subdisk	120
Viewing Volume Properties	122



Chapter 5. Volume Set Tasks	125
Volume Set Task Roadmap	125
VxVM Volume Sets	126
Displaying Volume Set Information	126
Creating a Volume Set	127
Adding a Volume to a Volume Set	127
Renaming a Volume Set	128
Recovering a Volume Set	128
Deleting a Volume Set	128
Removing a Volume from a Volume Set	129
Chapter 6. Using Volume Snapshots	131
Volume Snapshots Task Roadmap	132
Using the Volume Snapshot Features	133
Taking Snapshots of Volume Sets	133
Using the Instant Snapshot Feature	134
Using the FastResync Feature	134
Enabling FastResync	134
Creating Instant Snapshots	136
If FastResync is not Enabled	136
VxVM Selecting the Disk for FastResync	136
Selecting Your Own Disks for FastResync	139
Creating a Full-sized Instant Snapshot of a Volume	141
Creating a Volume for Snapshot	141
Creating Space-Optimized Snapshots of a Volume	143
Creating an Instant Snapshot of an ISP Volume	145
Creating an Instant Snapshot of an ISP Volume Set	146
Creating a Full-sized Instant Snapshot of a Volume Set	148
Creating Space-Optimized Snapshots of a Volume Set	148
Creating Break-Off Snapshots	150

Disabling FastResync	151
Refreshing an Instant Snapshot	151
Reattaching an Instant Snapshot	152
Restoring a Volume from an Instant Snapshot	152
Dissociating an Instant Snapshot	153
Splitting an Instant Snapshot Hierarchy	154
Start Synchronization	155
Adding Maps	155
Selecting Your Own Disks	155
VxVM Selecting the Disks	156
Using the Full Volume Snapshot Feature	156
Using Snapshot	157
Creating a Snapshot Copy of a Volume	157
Creating the New Volume	158
Merging a Snapshot Volume (Volume Snapback)	159
Dissociating a Snapshot Volume (Volume Snap Clear)	160
Aborting a Snapshot	160
Displaying the Snapshots for a Volume	161
Displaying Cache Volume Information	161
Backing up a Volume	162
Preparing to Restore a Volume From Backup	163
Chapter 7. File System Tasks	165
File System Task Roadmap	166
Creating a New File System on a Volume	167
Mount File System Details	169
Cluster Options	169
Permissions/Ownership	170
Advanced Options	170
Mounting a File System	171

Creating a New File System on a Volume Set	172
Unmounting a File System	172
Removing a File System from the File System Table	173
Defragmenting a File System	173
Creating a Snapshot Copy of a File System	174
Removing a Snapshot Copy of a File System	175
Checking a File System	176
Monitoring File System Capacity	177
Viewing File System Properties	177
Using the QuickLog Feature	178
Creating a QuickLog Volume	178
Enabling a QuickLog Device for a File System	179
Disabling a QuickLog Device for a File System	180
Removing a QuickLog Volume	181
Unmounting a File System from a Cluster Node	182
Removing Resource Information for a Cluster File System	183
Maintaining File Change Log	184
Enable File Change Log	184
Disable File Change Log	184
Synchronize File Change Log	184
Tune File Change Log	185
Maintaining Storage Checkpoints	186
Create a Checkpoint	186
Remove a Checkpoint	187
Mount a Checkpoint	187
Unmount a Checkpoint	188
Enable Checkpoint Quota	188
Disable Checkpoint Quota	189
Using Multi-Component File Systems	190
Setting up an MDS System	190

Mounting the Volume Set	190
Adding Volumes	190
Removing Volumes	191
Encapsulating Volumes	191
Set Intent Log Options	192
Chapter 8. Using the Quality of Storage Service Feature	193
QoS Task Roadmap	193
Configuring Relocation Policies	194
Viewing Relocation Policies and Modifying Policy Rules	194
Creating and Modifying Relocation Policies	197
Analyzing a File System	204
Relocating Files	207
Selecting Files Based on Maximum Age	207
Chapter 9. Device-Related Tasks	209
Device Task Roadmap	210
Displaying Lists of Supported Arrays	211
Enabling and Disabling Paths	212
Maintaining Disks	213
Maintaining the List of Foreign Devices	215
Display Information for a Foreign Device	215
Add Foreign Devices	215
Remove Foreign Devices	218
Chapter 10. Using Intelligent Storage Provisioning	221
Benefits of Using ISP	221
ISP Task Roadmap	222
Managing User Templates	223
Application Volumes	223
Intent	223

Template Set	224
User Template	224
Create a User Template	225
Create the User Template	225
Select Capabilities	226
Set Configuration Rules	227
Storage Selection Rules	227
Summary Page	229
Performance	229
Reliability	229
Edit a User Template	230
Delete a User Template	231
Managing Storage Pools	232
Data Pool	233
Clone Pools	233
Storage Pool Policy	233
Storage Pool Set	234
Creating a Storage Pool	235
Deleting a Storage Pool	236
Adding and Removing Storage Pool Disks	237
Renaming a Storage Pool	238
Associating a Template Set with Storage Pool	239
Disassociating a Template Set from a Storage Pool	240
Annotating a Disk	241
Storage Attributes	241
Launch the Annotate Service Dialog	242
Add a New User Attribute	243
Modify a User Attribute	244
Remove a User Attribute	244
Sort Disks by Attribute Name	245

Create a User Attribute Short-cut	246
Delete a User Attribute Short-cut	246
Chapter 11. Cluster Tasks	247
Cluster Task Roadmap	247
Working with Clusters	248
Creating a Shared Disk Group	249
Importing a Shared Disk Group	251
Maintaining Primary Policies	252
Setting the Primary Policy	252
Clearing the Primary Policy	253
Setting the Primary Node	253
Chapter 12. Troubleshooting	255
Alerts	255
Client Java Exceptions and Errors	255
Disk Troubleshooting	258
Disk State: Imported	258
Disk State: Free	258
Disk State: Not Setup	258
Disk State: Import Failed	259
Disk State: Disconnected	259
Disk State: Deported	259
Disk State: Offline	260
Disk State: Not Initialized	260
Disk State: Locked	260
Disk State: Inactive	261
Disk State: Foreign Owned	261
Volume Troubleshooting	262
Volume State: Stopped	262
Volume State: Uninitialized	262

Volume State: Degraded	263
Volume State: Failed	263
General Troubleshooting	264
Client/Server Problems	264
User is not authorized to connect to the server.	264
There is no such server.	264
Unable to connect to the server.	264
Possible client/server versions mismatch.	264
Server connection lost.	265
The server seems to have multiple addresses.	265
VEA version of plug-in component is not compatible. Cannot load plug-in .	265
Configuration Changes	265
Appendix A. Storage Selection Rules	267
Appendix B. Available ISP Definitions	269
Template Sets	270
ConfineVolume	270
DataMirroring	270
DataMirroringPrefabricatedRaid5	270
DataMirroringPrefabricatedStriping	270
DataMirrorStripe	271
DataStripeMirror	271
InstantSnapshottable	271
MultipathingThroughMirroring	272
MultipathingThroughMultiplePaths	272
PrefabricatedDataMirroring	272
PrefabricatedRaid5	272
PrefabricatedStriping	272
Raid5Templates	273
Striping	273




StripingPrefabricatedDataMirroring	273
Volume Templates	274
ArrayProductId	274
ColumnsOnSeparateComponents	274
ConcatVolumes	274
ConfineColumnsToSimilarStorage	274
ConfineLogsToSimilarStorage	274
ConfineMirrorsToSimilarStorage	274
ConfineToSimilarStorage	274
ConfineToSpecificStorage	274
DataMirroring	274
DataMirrorStripe	275
DataStripeMirror	275
DCOLogMirroring	275
DCOLogStriping	275
ExcludeSpecificStorage	275
InstantSnapshottable	275
LogsOnSeparateComponents	275
MirrorsOnSeparateComponents	276
MultipathingThroughMirroring	276
MultipathingThroughMultiplePaths	276
PrefabricatedDataMirroring	276
PrefabricatedRaid5	276
PrefabricatedStriping	276
Raid5LogStriping	276
Raid5Volume	276
Striping	277
Capabilities	278
ArrayProductId	278
ColumnsOnSeparateComponents	278



ConcatVolumes	278
ConfineColumnsToSimilarStorage	278
ConfineLogsToSimilarStorage	278
ConfineMirrorsToSimilarStorage	278
ConfineToSimilarStorage	279
ConfineToSpecificStorage	279
DataMirroring	279
DataMirrorStripe	279
DataRedundancy	279
DataStripeMirror	280
DCOLogMirroring	280
DCOLogStriping	280
ExcludeSpecificStorage	280
InstantSnapshottable	281
LogsOnSeparateComponents	281
MirrorsOnSeparateComponents	281
Multipathing	281
MultipathingThroughMultiplePaths	281
PrefabricatedDataMirroring	282
PrefabricatedDataRedundancy	282
PrefabricatedRaid5	282
PrefabricatedStriping	282
Raid5Capability	282
Raid5LogMirroring	282
Raid5LogStriping	283
Snapshottable	283
Striping	283
Storage Pools	284
any_volume_type	284
mirror_stripe_volumes	284

mirrored_prefab_raid5_volumes	284
mirrored_prefab_striped_volumes	285
mirrored_volumes	285
prefab_mirrored_volumes	285
prefab_raid5_volumes	285
prefab_striped_volumes	286
raid5_volumes	286
stripe_mirror_volumes	286
striped_prefab_mirrored_volumes	287
striped_volumes	287
Storage Pool Sets	288
mirrored_data_striped_clones	288
mirrored_prefab_raid5_data_mirrored_clones	288
mirrored_prefab_stripe_data_striped_clones	288
prefab_mirrored_data_prefab_striped_clones	288
stripe_mirrored_data_striped_clones	288
striped_prefab_mirrored_data_striped_clones	289
Glossary	291
Index	299





Preface

This guide describes how to use the VERITAS Volume Manager™ Graphical User Interface (GUI).

Audience

This guide is intended for system administrators responsible for configuring and maintaining systems using the VERITAS Volume Manager (VxVM) with the VERITAS Enterprise Administrator (VEA). This guide assumes the following:

- ◆ A basic understanding of storage management
- ◆ A basic understanding of system administration
- ◆ A working knowledge of the Solaris operating system
- ◆ Familiarity with VxVM and related concepts

Organization

This guide is organized as follows:

- ◆ Chapter 1, “Introducing the VEA GUI,” describes the main components of the GUI.
- ◆ Chapter 2, “Getting Started with VxVM VEA,” describes how to set up and start using VEA with VEA.
- ◆ Chapter 3, “Disk Tasks,” describes how to set up and use disks using VxVM with VEA.
- ◆ Chapter 4, “Volume Tasks,” describes how to create and use volumes and related objects using VEA with VEA.
- ◆ Chapter 5, “Volume Set Tasks,” describes how to use the volume sets feature.
- ◆ Chapter 6, “Using Volume Snapshots,” describes how to use the volume snapshot feature.
- ◆ Chapter 7, “File System Tasks,” describes the tasks associated with file systems.

Related Documents

- ◆ Chapter 8, “Using the Quality of Storage Service Feature,” describes how to use the Quality of Storage Service Feature.
- ◆ Chapter 9, “Device-Related Tasks,” provides information on the Device Discovery Layer feature.
- ◆ Chapter 10, “Using Intelligent Storage Provisioning,” provides information on using the VERITAS Intelligent Storage Provisioning (VERITAS Volume Manager Storage Administrator) feature.
- ◆ Chapter 11, “Cluster Tasks,” describes the optional cluster-related features.
- ◆ Chapter 12, “Troubleshooting,” provides information about VEA alerts, object states, and other troubleshooting topics.
- ◆ Appendix A, “Storage Selection Rules,” provides detailed ISP parameter information.
- ◆ Appendix B, “Available ISP Definitions,” provides ISP definitions.
- ◆ Glossary defines terms that relate to VEA.

Related Documents

The following documents provide related information:

- ◆ *VERITAS Volume Manager Administrator’s Guide*
- ◆ *VERITAS Volume Manager Troubleshooting Guide*
- ◆ *VERITAS Volume Manager Hardware Notes*
- ◆ *VERITAS Volume Manager Installation Guide*
- ◆ *VERITAS Volume Manager Release Notes*
- ◆ *VERITAS Intelligent Storage Provisioning Administrator’s Guide*
- ◆ VERITAS Volume Manager manual pages

The following table describes terms associated with the use of the mouse:

Term	Definition
<i>Click</i>	Press and release the mouse button.
<i>Double-click</i>	Click the left mouse button twice (quickly).
<i>Right-click</i>	Press and release the right mouse button.
<i>Press and Hold</i>	Press and continue to hold down the left mouse button.
<i>Point</i>	Move the cursor onto an item on the screen.

Term	Definition
<i>Select</i>	Click the left mouse button while the pointer is directly over the item to be selected.
<i>Drag</i>	Select the object and drag it while pressing the left mouse button.

Getting Help

If you have any comments or problems with the VERITAS products, contact VERITAS Technical Support:

- ◆ U.S. and Canadian Customers: 1-800-342-0652
- ◆ International Customers: +1 (650) 527-8555
- ◆ E-mail: support@veritas.com

For license information (U.S. and Canadian Customers):

- ◆ Phone: 1-650-527-0300
- ◆ E-mail: license@veritas.com
- ◆ Fax: 1-650-527-0952

For software updates:

- ◆ E-mail: swupdate@veritas.com

For information on purchasing VERITAS products:

- ◆ Phone 1-800-327-2232
- ◆ E-mail: sales.mail@veritas.com

For additional technical support information, such as TechNotes, product alerts, and hardware compatibility lists, visit the VERITAS Technical Support Web site at:

- ◆ <http://support.veritas.com> (U.S. and Canada Customers)

For additional information about VERITAS and VERITAS products, visit the Web site at:

- ◆ <http://www.veritas.com>

Downloading and Running VRTSexplorer

The `VRTSexplorer` program can help VERITAS Technical Support engineers diagnose the cause of technical problems associated with VERITAS products. You can download this program from the VERITAS FTP site or install it from the VERITAS Installation CD. For more information, consult the *VERITAS Volume Manager Release Notes* and the README file in the `support` directory on the VERITAS Installation CD.

Unique Message Number

If you encounter a product error message, record the unique message number preceding the text of the message. When contacting VERITAS Technical Support, either by telephone, or by visiting the VERITAS Technical Support website, be sure to provide the relevant message number. VERITAS Technical Support will use the message number to quickly determine if there are TechNotes or other information available to you.

A unique message number is an alpha-numeric string beginning with the letter "V." For example, in the message number:

```
V-5-732-8018 At least one disk must be specified
```

the "V" indicates that this is a VERITAS product error message. The text of the error message follows the unique message number.

Introducing the VEA GUI

1

Before using the VERITAS Volume Manager VEA GUI, you should familiarize yourself with the following introductory topics:

- ◆ VxVM Features
- ◆ Main Window
- ◆ Task Roadmap

VxVM Features

You use the VxVM to administer disks, volumes, and file systems on local or remote machines.

VxVM consists of a server and a client. The server runs on a UNIX machine that is running VxVM. The client runs on Windows, HP, Linux, Solaris, and AIX machines that support the Java Runtime Environment.

VxVM provides the following features:

- ◆ Accessibility - Using the Mouse and Keyboard

Accelerators and mnemonics are provided as alternatives to using the mouse (refer to the VEA online help for more information).
- ◆ Ease of Use

The task-based user interface provides access to tasks through menus or a task list. Administrators can easily navigate and configure their systems, and browse through all of the objects on the system or view detailed information about a specific object.
- ◆ Remote Administration

Administrators can perform VxVM administration remotely or locally. The VEA client runs on UNIX or Windows machines.
- ◆ Java-based Interface

A pure Java-based interface is used. Administrators can run VxVM as a Java application.

Main Window

- ◆ Scalability

VxVM can handle systems containing large numbers of disks. Administrators can view all of the objects on the system, or focus on a specific object or set of objects.
- ◆ Security

VxVM does not support communication between the client and server across a firewall.
- ◆ Multiple VEA Support

The VxVM client can provide simultaneous access to multiple host machines. The administrator can use a single client session to connect to multiple hosts, view the objects on each host, and perform administrative tasks on each host. Each host machine must be running the VxVM server.
- ◆ Multiple Views of Objects

VxVM provides multiple views of objects:

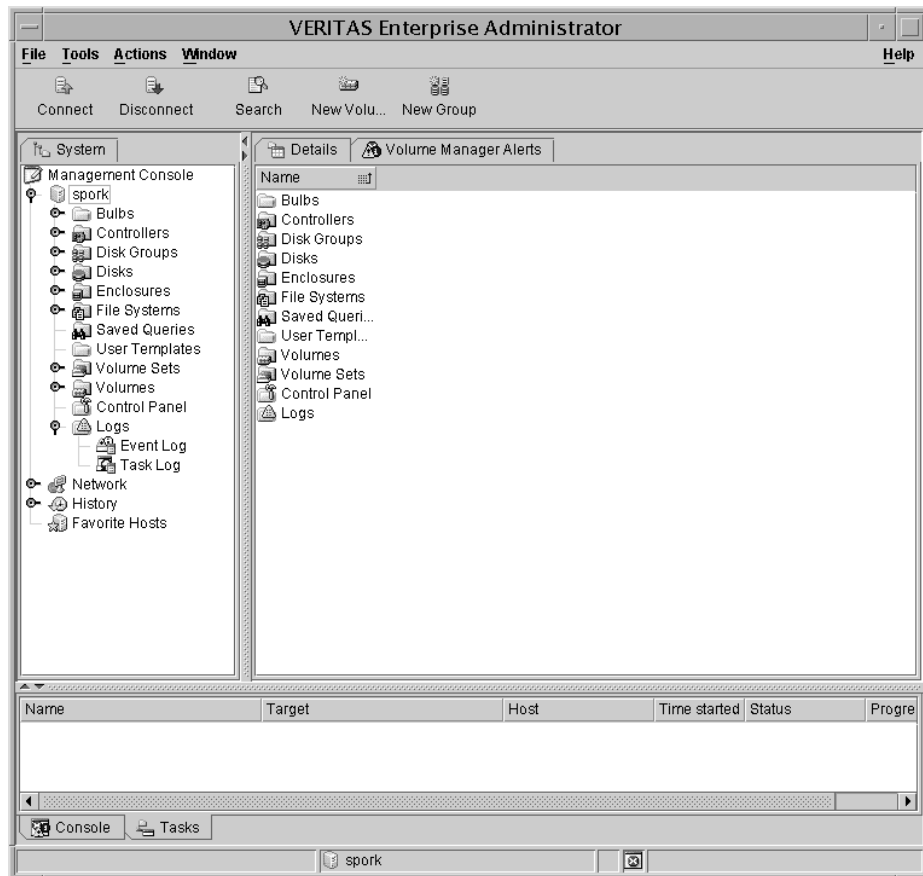
 - ◆ The main window tree and grid view lists objects on the system.
 - ◆ The Disk View window displays a graphical view of volumes, disks, and other objects in a disk group.
 - ◆ The Volume View window displays a graphical view of the volumes on a host.
 - ◆ The Volume to Disk Mapping window shows the relationships between volumes and their underlying disks.

Main Window

The VxVM Main Window consists of:

- ◆ The Object Tree and Grid
- ◆ The Menu Bar
- ◆ The Toolbar
- ◆ The Status Area

Main Window



The toolbar can be separated from and reattached to the main window (this is known as *docking*). The grid can be replicated in a separate window.

The Object Tree and Grid

The VxVM main window is divided into two main panes:

- ◆ Left Pane (Object Tree)

The following object groups typically appear as nodes in the object tree:

- ◆ **Hosts** – Any host machines connected to the current VxVM client session.
- ◆ **Controllers** – All controllers on the system.

Main Window

- ◆ **Disk Groups** – All disk groups on the system.
- ◆ **Disks** – All disks on the system.
- ◆ **Enclosures** – All enclosures (disk arrays) on the system.
- ◆ **File Systems** – All mounted file systems on the system.
- ◆ **Volumes** – All volumes on the system.
- ◆ **Right Pane (Grid)**

The right pane contains a *grid*, which is a tabular display of objects and their properties. The grid displays objects that belong to the group icon that is currently selected in the object tree (left pane).

The *splitter* is the vertical bar that separates the object tree from the grid. You can resize the left and right panes by pressing and holding the mouse button over the splitter and dragging the splitter to the left or right.

The Menu Bar

The menu bar at the top of the main window contains the following menus:

- ◆ **File**

The **File** menu establishes new host connections, disconnects current host sessions, provides access to an object Properties window, handles page setup, executes print and print preview functions, or exits VEA.
- ◆ **Tools**

The **Tools** menu allows you to:

 - ◆ Initiate the Setting Preferences dialog
 - ◆ Manage User Profiles
 - ◆ Maintain the VEA Error Console
 - ◆ Customize table headers (Customizing Table Headers)
 - ◆ Invoke the Set Up a Search Query feature
 - ◆ Initiate the Scan Disks function
- ◆ **Actions**

The **Actions** menu is a context-sensitive menu that launches tasks on a selected object. It is inactive until you connect to a host. The **Actions** menu is dynamic and changes its options based on the type of object that is selected. When an object is selected, the **Actions** menu provides access to tasks appropriate for the selected object.
- ◆ **Window**

The **Window** menu gives you access to the Tear off feature. This allows you to view several different object types at the same time, and to correlate relationships between them.

Window -> Tear off

- ◆ **Help**

The **Help** menu provides access to online help for VEA.

Customizing Table Headers

You can use the Tools menu to customize table headers, that is, reorder, remove, or add, any of the table headings and associated columns.

▼ To customize table headers

1. Choose **Tools > Customize Table Headers**. The Customize Columns dialog box is displayed.
2. Select the table to be modified from the **Type:** drop-down menu. The columns applicable to that table are then displayed under **Show columns:** (those columns that are currently shown) and **Free columns:** (those columns that are currently hidden).
3. You can now customize your tables by reordering (**Move Up, Move Down**), displaying (**Add**), and hiding (**Remove**) headers and columns.
4. Click **OK** to complete the changes.

The Toolbar

The toolbar provides quick access for performing some of the program's most common tasks.

The *toolbar handle* is the thin bar next to the toolbar. You can use the toolbar handle to separate the toolbar from the main window or move the toolbar to the bottom, side, or top of the main window. To reposition the toolbar, press and hold the mouse button over the toolbar handle and drag the toolbar to its new location.

Note The toolbar handle is not visible when running on UNIX clients.

The Status Area

The status area is at the bottom of the main window. When an object fails or experiences an error, an alert (error) icon appears in the status area. The Alert Monitor window provides details about the error. You can access the Alert Monitor window by double-clicking on the alert icon in the status area.

Task Roadmap

VxVM tasks are grouped into:

- ◆ Setup Tasks
- ◆ Maintenance Tasks
- ◆ Repair/Recovery Tasks

Setup Tasks

If you are using VxVM and VEA for the first time, you need to install both products, place your disks under VxVM control, and create volumes.

▼ Setting Up Your System

- ◆ Install and initializing VxVM.
- ◆ Install VEA.
- ◆ Start VEA.

These tasks are described in the *VERITAS Volume Manager Installation Guide*.

▼ To set up disks

- ◆ Place disks under VxVM control.
- ◆ Create additional disk groups for the new disks (optional).
- ◆ Designate one or more disks as hot-relocation spares (optional).

These tasks are described in Chapter 3, "Disk Tasks."

▼ To set up volumes

- ◆ Create volumes.
- ◆ Create mirrors for any volumes that need to be redundant.

These tasks are described in Chapter 4, "Volume Tasks."

▼ To set up file systems

- ◆ Placing file systems on volumes.
- ◆ Mounting a file system.

These tasks are described in Chapter 7, “File System Tasks”.

▼ To set up clusters

- ◆ Set up a shared disk group.

These tasks are described in Chapter 11, “Cluster Tasks.”

Maintenance Tasks

After you have set up your disks and volumes, you need to perform some or all of the following maintenance tasks.

▼ Monitor objects

- ◆ Use the Main Window tree (left pane) and grid (right pane) to view disks, volumes, file systems, and other objects on the system.
- ◆ Use the Disk View Window to display a graphical view of disks, volumes, and other objects in a particular disk group.
- ◆ Use the Volume View window to display a graphical view of the volumes for a host.
- ◆ Use the Volume to Disk Mapping Window to view the relationships between volumes and their underlying disks.
- ◆ Use the Object Properties Window to view detailed information about a selected object.
- ◆ Use the Alert Tab to view information about errors or failures associated with objects that are marked with alert icons.

(Windows are described in Chapter 2, “Getting Started with VxVM.”)

▼ Maintain disks

- ◆ Adding a Disk to VxVM.
- ◆ Creating a Disk Group.
- ◆ Upgrading a Disk Group Version.

These tasks are described in Chapter 3, “Disk Tasks.”

Task Roadmap

▼ **Maintain volumes**

- ◆ Renaming a Volume as users and applications require more space.
- ◆ Changing the VxVM ISP Column Count.
- ◆ Adding a Mirror to a Volume to any unmirrored volume that should be redundant.

These tasks are described in Chapter 4, "Volume Tasks."

▼ **Maintain file systems**

- ◆ Creating a New File System on a Volume Set.
- ◆ Increase the size of file systems as users and applications require more space.
- ◆ Create snapshot copies of file systems and then back up the file systems.

These tasks are described in Chapter 7, "File System Tasks".

▼ **Maintain clusters**

- ◆ Creating a Shared Disk Group
- ◆ Importing a Shared Disk Group

These tasks are described in Chapter 11, "Cluster Tasks."

Repair/Recovery Tasks

If disk and/or system failures occur, VxVM generally performs any necessary recovery operations automatically. However, depending on the circumstances, you may need to attempt to recover from the situation as follows.

▼ **Protect volumes from disk failure**

- ◆ Adding a Disk to VxVM.
- ◆ Replacing a Disk.

These tasks are described in Chapter 3, "Disk Tasks."

▼ **Repair/recover volumes**


- ◆ Recovering a Volume.
- ◆ Recover all volumes on a failed disk.
- ◆ Recovering Volumes in Disk Groups.
- ◆ If a volume is not recoverable, restore the volume from backup.

These tasks are described in Chapter 3, "Disk Tasks," and Chapter 4, "Volume Tasks."



▼ **Repair/recover file systems**

- ◆ Perform a file system consistency check on any file systems that may be damaged.
(Refer to Chapter 5, "File System Tasks.")





Task Roadmap



Getting Started with VxVM VEA

2

Before using VEA, you need to familiarize yourself with the following:

- ◆ Setting Up Your System
- ◆ Starting VEA
- ◆ Viewing Objects and Object Properties
 - ◆ Tree and Grid View
 - ◆ Displaying Disk Information
 - ◆ Volume View Window
 - ◆ Displaying Volume Information
 - ◆ Volume View Window
 - ◆ Volume Layout Details Window
 - ◆ Volume to Disk Mapping Window
- ◆ Viewing Task History
- ◆ Dissociated Plex View Window
- ◆ Command Log
- ◆ Using the Search Feature
- ◆ Edit and Remove Search Queries
- ◆ Using the Search Feature
- ◆ Performance Tuning
- ◆ Load Balancing
- ◆ Setting Preferences
- ◆ Disk Failure or Removal

Setting Up Your System

VEA consists of a server (`vxsvc`) and a client (`vea`). You must install and run the VEA server on a UNIX machine running VxVM. You can install and run the VEA client on any Windows, HP, Linux, Solaris, or AIX machine that supports the Java Runtime Environment.

Before you can use VEA, you need to:

1. Install VxVM on the machine to be administered.
2. Install the VEA server and client on the machine to be administered.
3. If users other than root need to access VEA, set up security to specify which users can run VEA.
4. If you plan to run the VEA client from a machine other than the machine to be administered, install the VEA client on the machine where the client will run.
5. Start the VEA server on the machine to be administered.
6. Start the VEA client.

Refer to the *VERITAS Volume Manager Installation Guide* for details on how to install and set up VEA.

Starting VEA

Only users with appropriate privileges can run VEA. (Refer to *VERITAS Volume Manager Installatin Guide* for information on security and access permissions.)

You can use VEA to administer the local machine or a remote machine. The VxVM daemon (`vxconfigd`) and the VEA server (`vxsvc`) must be running on the machine to be administered. Once VEA is installed, the VEA server starts up automatically at system startup.

Viewing Objects and Object Properties

VEA provides the following views of objects and their properties:

- ◆ Tree and Grid View

The main window tree and grid view displays volumes, disks, file systems, and other objects on the system. The grid also displays a limited set of properties for listed objects.

- ◆ Disk View Window

The Disk View window displays a graphical view of disks, volumes, and other objects in a particular disk group.

- ◆ Volume View Window

The Volume Layout Details window displays a close-up graphical view of a single volume and its components (mirrors/plexes and subdisks).

- ◆ Volume to Disk Mapping Window

The Volume to Disk Mapping window shows the relationships between volumes and their underlying disks.

- ◆ Object Properties Window

The object Properties window contains detailed information about a specific object.

Tree and Grid View

The main window tree and grid view displays volumes, disks, file systems, and other objects on the system. The object tree is a hierarchical display of object groups. The grid is a tabular display of objects and a limited set of properties. VEA constantly monitors objects on the system and makes appropriate changes to the tree and grid displays.

Disk Usage Legend

The disk usage legend is displayed at the bottom of the right pane; and shows how to interpret the colors with respect to Disk Usage (Normal Usage, High Usage, and Critical Usage).

Displaying Disk Information

Click on a disk to display information about that disk. An appropriate subset of tabs is also provided from the following:

- ◆ Disk Attributes
- ◆ Disk View
- ◆ Volume Manager Alerts
- ◆ Disk Regions

Viewing Objects and Object Properties


- ◆ Volumes
- ◆ Controllers
- ◆ Paths

Disk View Window

The Disk View window displays a graphical view of disks. It is dynamic, so the objects displayed in this window are automatically updated when their properties change. You can select objects or perform tasks on objects in the Disk View window.

The Disk View window can be accessed by clicking on a disk and selecting **Actions > Disk View**.

Disk View Window



You can move the horizontal splitter to adjust the relative heights of both panes.

Using the Menu Bar

The menu bar at the top of the Disk View window contains the following menus:

- ◆ **File**

The **File** menu enables you to:

- ◆ Open the disk group in which the selected disk resides
- ◆ Access to the page setup menu
- ◆ Print the Disk View
- ◆ Close the Disk View window

- ◆ Access the Properties window
- ◆ Display selected disks
- ◆ **Options**

The **Options** menu provides you with the following selections:

 - ◆ Refresh View
 - ◆ Expand/Collapse—provides detailed information for all disks in the view
 - ◆ Vol Details—basic volume information
 - ◆ Full Volume Display—volume objects also display the volume size, free space, and state of the volume
 - ◆ Projection—highlights associated objects

With the exception of Refresh View and Full Volume Display, these options are also available from the icons under the Disk View menu bar.

- ◆ **Action**

The **Action** menu provides all the tasks available on the corresponding object in the main window.

Using the Toolbar

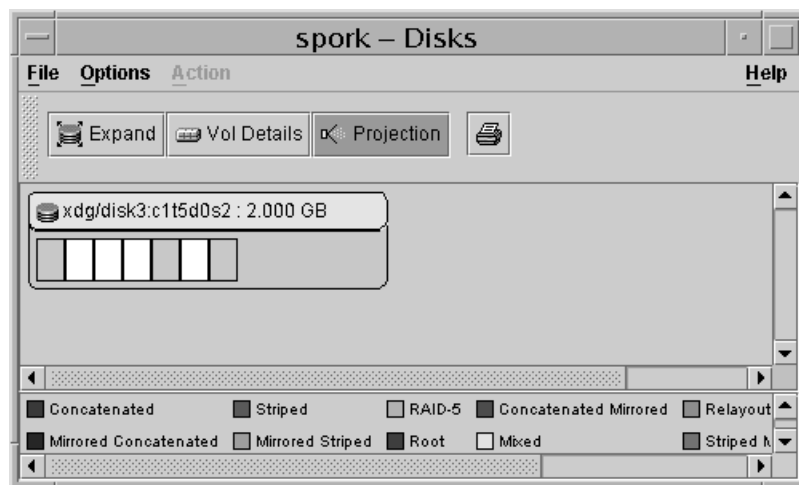
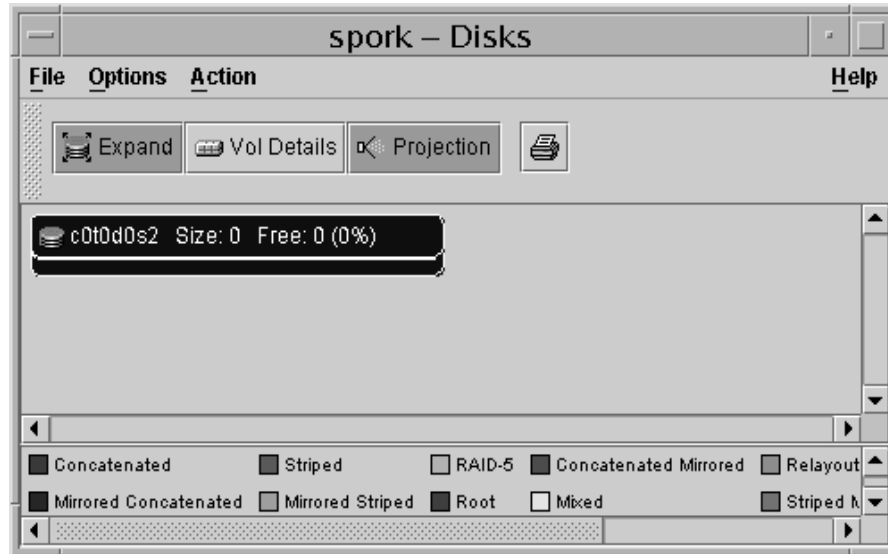
The toolbar at the top of the Disk View window has the following set of buttons:

- ◆ **Expand**

The **Expand** button shows more detailed information about all disks in the Disk View, including subdisks and free space. This button toggles back and forth between objects in layout mode or detailed mode, depending on the previous mode.

Using the Expand button

Viewing Objects and Object Properties



- ◆ **Projection**

The **Projection** button highlights objects associated with the selected subdisk or volume.

Projection shows the relationships between objects by highlighting objects that are related to or part of a specific object. When disks are in layout mode or detailed mode, clicking on a subdisk highlights the location of the subdisk on the corresponding disk.

◆ **Print**

The **Print** button prints the Disk View.

The Disk View window displays plexes and subdisks that are associated with volumes. To view Volume Manager objects that are not associated with volumes (such as dissociated plexes), open the Dissociated Object View window (see **Dissociated Plex View Window**).

Moving Disks in the Disk View

In the Disk View, you can move subdisks by dragging the subdisk icons to their new locations. You can drag a subdisk to another disk or a gap on the same disk. To move a subdisk, place the pointer over the subdisk and then press and hold the mouse button to drag the subdisk to another location.

Dragging and Dropping of subdisks can be enabled and disabled (see Setting Preferences for details).

Caution Moving subdisks reorganizes volume disk space and must be done with care.

Mapping Disks to Volumes

You can map a disk to the associated volume by bringing up both the Disk View and the Volume View (see Volume View Window). When you click on the disk in the Disk View, the corresponding plex in the Volume View is highlighted.

Displaying Volume Information

Click on a volume to display information about that volume. An appropriate subset of tabs is also provided from the following:

- ◆ Data Change Object (DCO)
- ◆ Disk View
- ◆ Volume Manager Alerts
- ◆ Disks
- ◆ Mirrors
- ◆ Subdisks
- ◆ Snapshot Volumes
- ◆ Logs

Viewing Objects and Object Properties

Volume View Window

The Volume View window gives the characteristics of the volumes on the disks.

The Volume View window can be accessed by clicking on a volume and selecting **Actions > Volume View**.

The menu bar at the top of the Volume View window contains the following menus:

- ◆ **File**

The **File** menu prints the properties of a selected object, closes the Volume View window, and provides access to the Properties window.

- ◆ **Action**

The **Action** menu provides access to all the tasks available on the corresponding object in the main window.

The toolbar at the top of the Volume View window has the following set of buttons:

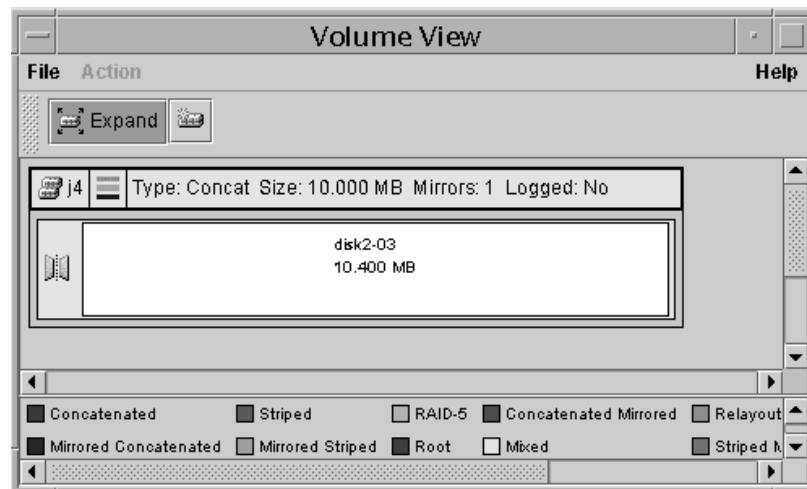
- ◆ **Expand all volumes**

The **Expand Volume** button shows more detailed information about all disks in the Volume View. This button displays objects in layout mode or detailed mode, depending on the previous mode.

- ◆ **New volume**

The **New Volume** button brings up the create volume dialog box.

Using the Create Volume Dialog Box

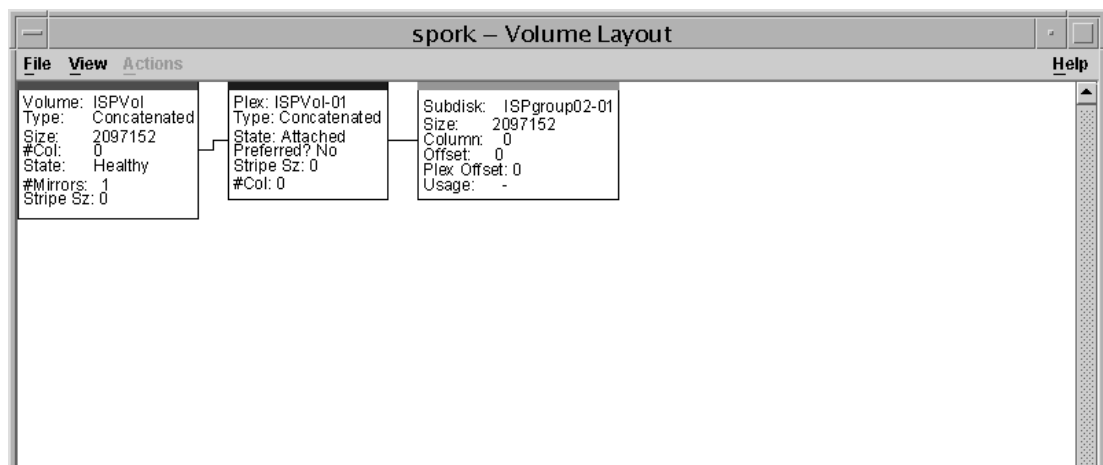


Volume Layout Details Window

The Volume Layout Details window displays a graphical view of the selected volume's layout, components, and properties. You can select objects or perform tasks on objects in the Volume Layout Details window. The Volume Layout Details window is dynamic, so the objects displayed in this window are automatically updated when the volume's properties change.

The **View** menu changes the way objects are displayed in this window. To access the popup menu for an object, right-click on the object, or use the **Actions** menu.

Volume Layout Details Window



You can view objects in the Volume Layout Details window in the following ways:

- ◆ To display the Volume Layout Details window for a volume, click on the volume (in the main window grid) and then choose **Actions > Layout View**.
- ◆ To view a different volume, go back to the main window and select the required volume.
- ◆ To access the popup menu for an object, right-click on the object.

Use the **File** menu to close the window.

Volume to Disk Mapping Window

The Volume to Disk Mapping Window displays a tabular view of volumes and their underlying disks. Volumes are listed in the top row of the table and disks are listed in the left column of the table. Each circle icon in the table indicates that part of the corresponding volume is located on the corresponding disk.

Viewing Objects and Object Properties

This window can also display details such as the subdisks on each disk. Each subdisk or gap is listed with its disk offset and length. When subdisks are shown, circles show the relationship between volumes and subdisks.

The Volume to Disk Mapping window is dynamic, so the contents of this window are automatically updated when objects are added, removed, or changed. For instructions on navigating the Volume to Disk Mapping window, see “Navigating the Volume to Disk Mapping Window” on page 21.

- ◆ To display the Volume to Disk Mapping window, click on the disk group (in the main window grid) and then choose **Actions > Disk/Volume Map**.

The **File** menu displays volume and disk property information, or closes the window.

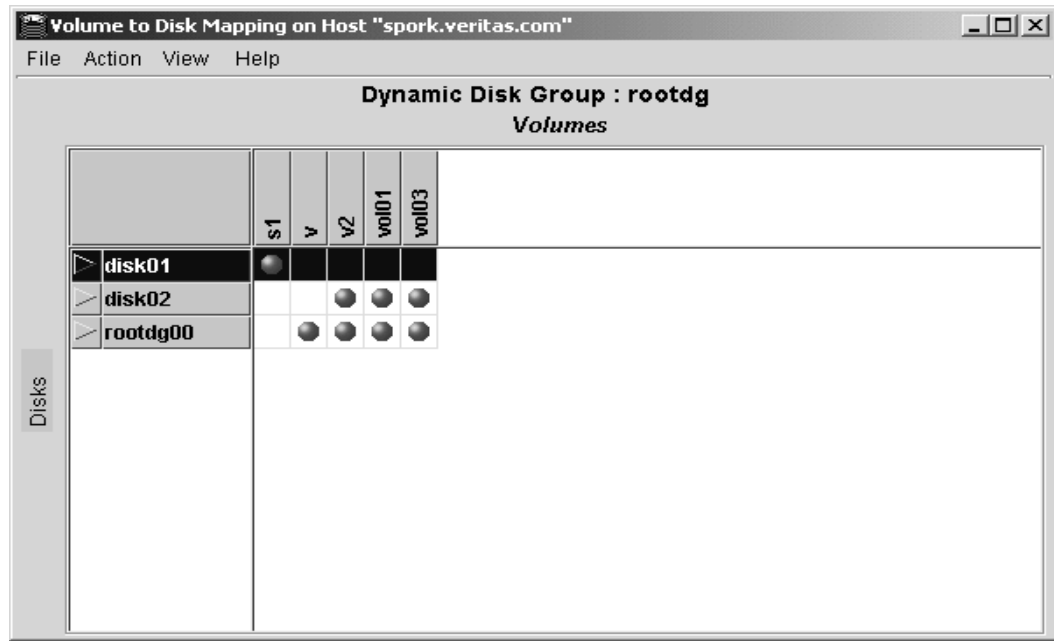
The **View** menu changes the way objects appear in the Volume to Disk Mapping window, and allows you to set up collection of statistics for the disk group.

The context-sensitive **Actions** menu allows you to perform tasks for a selected object. To access the popup menu for an object, right-click on the object.

The Volume to Disk Mapping window has a performance monitoring feature that ranks volume response time. When performance monitoring is turned on, the icons in the table change. The volumes with the slowest reads/writes are shown as red icons. Volumes with relatively slow response time are shown as yellow icons. All other volumes are shown as green icons. For instructions on viewing real-time statistics, see “Viewing Real-Time Statistics” on page 34.

“Volume to Disk Mapping Window” shows the relationships between volumes and disks.

Volume to Disk Mapping Window



Navigating the Volume to Disk Mapping Window

You can view volumes and their associated disks using the following procedures:

- ◆ To display the Volume to Disk Mapping window for all volumes in a disk group, click on the particular disk group you want to view under the Disk Groups tree (in the main window grid) and then choose **Actions > Disk/Volume Map**.
- ◆ To display all of the subdisks on a particular disk, click the arrow to the left of the disk name. Any subdisks are listed below the disk name. To hide the subdisks, click the arrow again.
- ◆ To display all of the subdisks on all disks, choose **View > Expand All**. To hide all of the subdisks on all disks, choose **View > Collapse All**.
- ◆ To use the statistics collection feature, see Viewing Real-Time Statistics.
- ◆ To highlight a disk row, click on the disk name. To highlight a volume column, click on the volume name. To highlight both the disk and volume associated with a particular circle, click on the circle.
- ◆ Volumes with relatively slow read/write performance are shown as circular red icons.
- ◆ To print the map then choose **File > Print** and complete the Print dialog box.

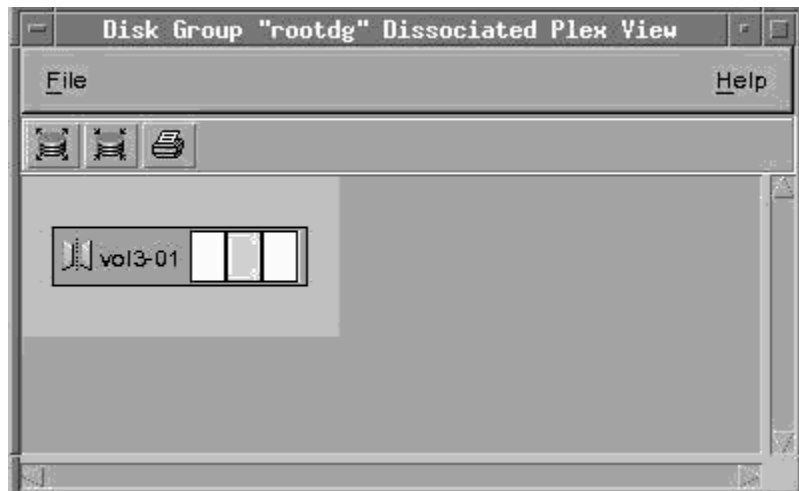
Mapping Volumes to Disks

You can map a plex to the associated disk by bringing up both the Volume View and the Disk View. When you click on the plex in the Volume View, the corresponding disk in the Disk View is highlighted.

Dissociated Plex View Window

The Dissociated Plex View window shows the dissociated plexes for a particular disk group.

Dissociated Plex View Window



The Dissociated Plex View window can be displayed by clicking on a disk group and selecting **Actions > Dissociated Plex View**.

Object Properties Window

The Object Properties window contains detailed information about the selected object. The Properties window contains a set of tabbed pages, each of which contains information about the object and related objects. The tab labels and page contents vary, depending on the type of object selected. You can display a different page by clicking on the appropriate tab label. Some object properties can be changed through the Properties window for the object. For instructions on navigating the Object Properties window, see “Navigating the Object Properties Window” on page 23.

“Volume Properties Window” shows how a volume appears in the Volume Properties window.

Volume Properties Window



Navigating the Object Properties Window

You can view object properties using the following procedures:

- ◆ To view the properties of an object, click on the object (in the main window grid) and then choose **Properties** from the popup menu.

Example:

To display volume properties, click on a volume and then choose **Volumes > Properties**.

- ◆ To go to a tab in a properties window, click the appropriate tab label.
- ◆ To change any of the editable items in the Properties window, make the appropriate changes and then click **OK**. This changes the settings for *all* properties tabs in the Properties window.

Multiple Selections

Multiple Selections

You can select more than one object from a list (for example, when you need to add several disks into a disk group) using the SHIFT and CONTROL keys in conjunction with the left mouse button.

Use the SHIFT key to select a range of objects.

Use the CONTROL key to select individual objects.

Accessing Tasks

Most VEA tasks are performed by selecting objects and/or tasks and then completing the resulting dialog boxes. With VEA, you can access tasks from the menu bar and from a context-sensitive popup menu.

Viewing Task History

VEA logs all task requests. You can view the task history (including tasks in progress) using the Task History window, or the Command Log file.

The Task History window displays a history of tasks performed in the current session (and any other sessions running on the machine). Each task is listed with properties such as task name, task status, and start time.

To display the Task History window, click on **Tasks** at the bottom of the main window.

Task History Window



For a selected task, you can use the popup menu options: **Abort Task** to cancel the task; **Pause Task** to temporarily stop a task, **Resume Task** to continue the task; and **Throttle Task** to throttle a task.

Task Throttling

Task Throttling causes an operation to pause for the allowed amount of time whenever a disk I/O is performed, allowing the CPU and disks to perform other tasks.

Select the popup menu option **Throttle Task** to change the task priority. You can change the priority for a single task, or select **Throttle All Tasks** to apply the priority to all Volume Manager Tasks. Click **Ok** to complete your selection.

Clearing Finished Tasks

Tasks are persistent in the Tasks console. To remove completed tasks from the window, select **Clear All Finished Tasks** from the popup menu.

Viewing Task Properties

To view task properties and the low-level commands used to perform a task, select the task and choose **Properties** from the popup menu.

Command Log

The command log file (`veacmdlog`) contains a history of VEA, VxVM, and VxFS tasks performed in the current session and previous sessions. It contains a description of each task, along with properties such as the start/finish times and the low-level commands used to perform the task. For failed tasks, the command log includes any relevant error messages. The command log file resides, and can be viewed, at `/var/adm/vx/veacmdlog`.

The following example shows a sample command log file entry:

```
Description: Add Log to Volume
Date: Wed April 17 4:36:11 2002
Command: /usr/sbin/vxassist -g datadg addlog data vol
Output: vxvm:vxassist: ERROR: Volume datavol-L01: Volume's logging
type not supported
Exit Code: 20
```

Using the Search Feature

The Search feature allows you to search and display information for objects on the system running VxVM. You can:

- ◆ Set Up a Search Query
- ◆ Use a Saved Query to Search and Display Information
- ◆ View the Results of a Previously Saved Query
- ◆ Edit and Remove Search Queries

Set Up a Search Query

The Search feature allows you to search and display information for objects on the system running VxVM.

▼ **To set up a new query**

1. Choose **Tools > Search**. This brings up the “Search” dialog.
2. Click **New Query**.
3. In order to find an object, you must first select the object type (under **Look for:**).

Object types include:

- ◆ Volumes
- ◆ Volume Sets
- ◆ Disks
- ◆ Disk Groups
- ◆ Controllers
- ◆ Enclosures
- ◆ File systems

4. Once you have chosen the object type, you define the search criteria (under **Define Criteria:**):

The criteria menu selection is based on the object type you have selected.

If you chose Volumes, the available Attributes are:

- ◆ Name
- ◆ Disk group name
- ◆ Size
- ◆ Status
- ◆ Disk
- ◆ Device
- ◆ # Mirrors
- ◆ Layout
- ◆ Usage

If you chose Volumes sets, the available Attributes are:

- ◆ Name
- ◆ Disk group name

If you chose Disks, the available Attributes are:

Using the Search Feature

- ◆ Name
- ◆ Disk group name
- ◆ Controller name
- ◆ Device name
- ◆ Status
- ◆ Size
- ◆ Free Space

If you chose Disk Group, the available Attributes are:

- ◆ Name
- ◆ Imported
- ◆ Size
- ◆ Free space
- ◆ # Volumes
- ◆ # Disks

If you chose Controllers (only shown if DMP is installed), the available Attributes are:

- ◆ Physical name
- ◆ Logical name

If you chose Enclosures (only shown if DMP is installed), the available Attributes are:

- ◆ Name
- ◆ Type
- ◆ Connected

If you chose File Systems, the available Attributes are:

- ◆ Mount point
- ◆ Type (vxfs, ufs, hfs, ext2, jfs)
- ◆ Free space
- ◆ Capacity

5. Once you have chosen an Attribute, you must choose how to compare the object.

The possible ways to compare the object change with your previous selections, and are limited to:

- ◆ Any

- ◆ Contains
 - ◆ At least
 - ◆ At most
 - ◆ Is
 - ◆ Is not
 - ◆ True
 - ◆ False
6. If you did not select the Usage, Status or Imported attribute, you must enter a number into the value box. The attribute of the object type you selected will be compared with the value you enter in the way you selected. If you selected the Status attribute, you will need to select a value from Uninitialized, Healthy, Stopped, Degraded, and Failed.
 7. If you selected the Size or Free space attribute, you will also need to define the unit of measurement (Sectors, KB, MB, GB, TB).
 8. If you selected the Usage attribute, you will also need to define the type: Cache, Data Change Object, Volume Set Component, Volume, or Data (regular volume).
 9. Use the **Add** and **Remove** buttons to add and remove search statements.

Note You can construct an 'or' statement for exact matches ('Is' criteria) by clicking **Add** after entering a value; then entering a new value and again clicking **Add**. Similarly, you can construct an 'and' statement for 'Is not' and 'contains' criteria.

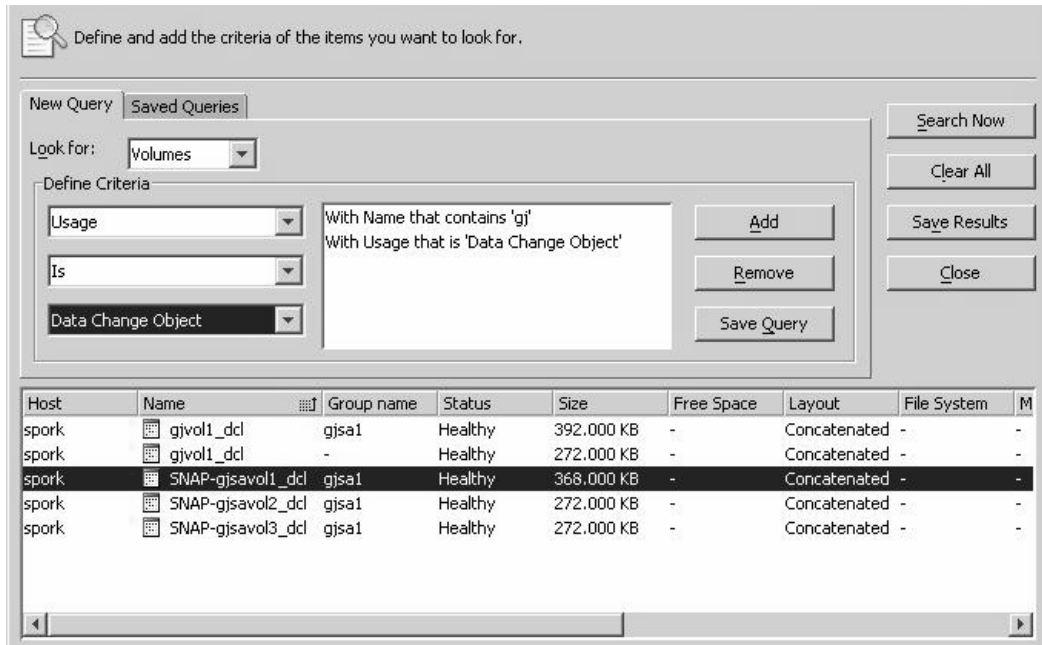
10. If you wish to save the query data, click the **Save Query** button.
 - ◆ The Save Query dialog box appears. Select the name for the saved query (**Name:**), and, if you want, supply a description of the query (**Description:**). You can display the saved queries in the Main Window by selecting the **Active** checkbox.
 - ◆ Click **OK** to save the query and close the dialog box.
11. Click the **Search Now** button to perform the search.

Example:

The following screen shows the results for a search for a Volume (**object type**), with Name (**attribute**) containing "gj," and Usage (**attribute**) for "Data Change Object.

Using the Search Feature

Volume Search



12. If you wish to save the search results, click the **Save Results** button.

- ◆ The Save Dialog box appears. Select the path or folder you want the file to save to.
- ◆ Enter a name for the file into the **File Name:** field.
- ◆ Click **Save** to save the file and close the dialog box.

Use a Saved Query to Search and Display Information

You can perform a search using a previously saved query.

▼ To search and display information

1. Choose **Tools > Search**. This brings up the Search dialog.
2. Click **Saved Queries**.
3. Select an already saved query from the **Query** drop-down list.
4. Click **Search Now**. The information is displayed at the bottom of the dialog.
5. Click **Close** when you have finished.

View the Results of a Previously Saved Query

Queries are saved in the **Saved Queries** node in the left pane. You can display the information from a previously saved query by expanding the appropriate query.

Edit and Remove Search Queries

You can edit a saved query description, or remove existing queries.

▼ To edit a saved query

1. Select the query to be edited from the **Saved Queries** node, then choose **Actions > Edit**. This brings up the **Edit Query** dialog.
2. You can now edit the name and description of the query.
3. Select the **Active** if you want to display the results in the main window.
4. Click **OK**.

▼ To remove a query

1. Select the query to be removed from the **Queries**, then choose **Actions > Remove**. This brings up the **Remove Query** dialog.
2. Click **Yes** to confirm removal.

Performance Tuning

When a program or device uses too much disk I/O, bottlenecks can occur, creating a hot spot and degrading performance. You can make adjustments by moving subdisks from an area of high I/O usage to another disk that has lower usage.

The VxVM statistics function provides I/O statistics to allow you to do performance tuning to improve overall disk and system performance.

You can improve disk and system performance by identifying I/O bottlenecks, or *hot spots*; then moving data among physical drives to evenly balance the I/O load among the disk drives. This is known as *load balancing*.

See also:

- ◆ Real-time Statistics Collection and Display
- ◆ Load Balancing
- ◆ Disk Failure or Removal
- ◆ Disk Phaseout

Real-time Statistics Collection and Display

To set up real-time statistics collection and display, you need to Select Data Gathering Options.

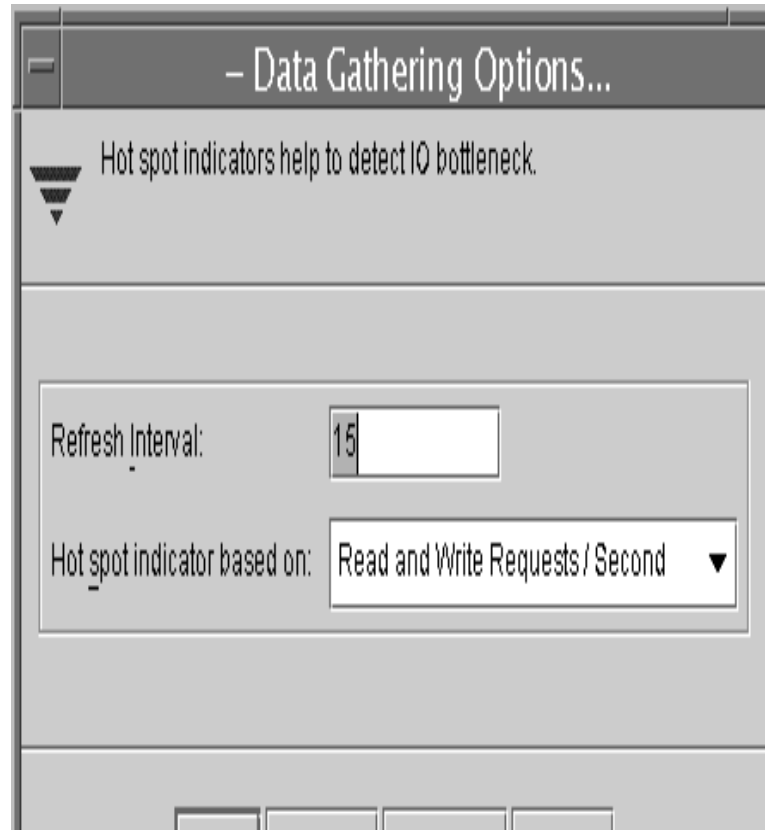
Select Data Gathering Options

Use the following procedure to set the display options.

▼ To select the online data gathering options:

1. Select a disk group to be monitored.
2. Choose **Actions > Statistics View>** to bring up the Volume to Disk Mapping window.
3. Choose **View > Data Gathering Options...** to bring up the data gathering options dialog box.

Data Gathering Options dialog box



4. Select the pulldown menu for the **Hot spot indicator based on** categories.

Note Each time you select an indicator and “Set Refresh Interval,” , you must click **OK** and move onto the next category you wish to set up by Choose **View > Data Gathering Options...**

- ◆ **Read Requests/Second**
The number of read requests per second for selected storage objects.
- ◆ **Write Requests/Second**
The number of write requests per second for selected storage objects.
- ◆ **Read Blocks/Second**
The amount of read request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the read requests made.

Performance Tuning

- ◆ **Write Blocks/Second**

The amount of write request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the write requests made.

- ◆ **Avg Time (microseconds)/Read Block**

The average time in microseconds that it takes to process a read block from the time a request is made until the data is returned.

- ◆ **Avg Time (microseconds)/Write Block**

The average time in microseconds that it takes to process a write block from the time a request is made until the write status is returned.

- ◆ **Read and Write Requests/Second**

The number of read and write requests per second for selected storage objects

- ◆ **Read and Write Blocks/Second**

The amount of read and write request data (in blocks per second) that is processed for selected storage objects. It is the throughput on the read and write requests made.

By default, all the statistical parameters will be enabled.

5. Set **Refresh Interval**

This sets the time interval (in seconds) for refreshing the data collection screen display. The program continuously runs the actual statistics in the background, but changes their display on the Online Monitoring according to the refresh rate you set. The default threshold is 15. The threshold is the minimum amount of I/O per second that a disk or subdisk has to have in order to be categorized with an I/O activity higher than average.

6. When you are finished with your changes, click **OK**.

Viewing Real-Time Statistics

You can view real-time statistics in the Volume to Disk Mapping Window. By default, the Volume to Disk Mapping Window comes up with statistics collection switched off. To turn it on, you need to:

- ◆ Select the disk group to be monitored from VEA.
- ◆ Select **Actions > Disk/Volume Map**
- ◆ From the Volume to Disk Mapping window, select:
 - ◆ **View > Collect Statistics**

Statistics are accumulated from when they were reset (usually they are reset when the system boots/reboots, or at object creation, whichever is most recent).

- ◆ **View > Compare on Refresh**

Each statistic request is compared with the last one to determine where the current high I/O activity is occurring.

Load Balancing

If disk activities are heavily concentrated on one or a small number of disks in the storage subsystem, it may create bottlenecks. You can use the **Moving a Subdisk** and possibly the **Splitting a Subdisk** features to spread out disk accesses more evenly across all the disks to balance the load.

If a disk has High or Critical I/O activity (shown by a red or yellow pie symbol), you may consider moving one or more of its subdisks to another disk that shows below average I/O activity (shown by a blue pie symbol). The idea is to move just enough activity to achieve balance. A careful study of the statistics for the disk with Critical activity may identify the best subdisks to move. You should move subdisks only when a disk has High or Critical I/O activity over a prolonged period of time and performance is affected. Moving a subdisk to another disk has an effect on I/O as well, but it should be compensated for by the other disk having much lower I/O activity. You would need to look at the statistics after the subdisk move to see whether the move was effective in balancing the load.

Disk Failure or Removal

Moving a subdisk for redundant volumes (mirrored or RAID-5) will use the redundant data to recreate the subdisk on the healthy disk. However, for nonredundant volumes (concatenated or striped), the data cannot be recreated and doing subdisk move will therefore lose data, which could be recovered if the disk can be repaired. Thus, when you attempt to move a subdisk from a failed or missing disk that has nonredundant data, a dialog box comes up that asks you if you want to force the move. You may want to force the move if you do not need the data anymore or you can no longer recover the data. By doing so, you will retain the volume structure but there is no guarantee that the data will be recoverable.

Use the **“Moving a Subdisk”** on page 117 feature to move the part of a volume that is on a failed or missing disk to a healthy one.

Setting Preferences

Disk Phaseout

When a disk starts getting intermittent I/O errors and shows signs of hardware fatigue, you can use the **Moving a Subdisk** feature to move all its subdisks to healthier disks. The benefit of moving subdisks instead of copying the volumes is that you need only enough space on the receiving disks for the subdisks on the one failed disk, not for entire volumes that may span multiple disks. Another advantage is that there is no interruption in I/O.

Moving subdisks provides you a flexible means of making adjustments in your storage system while it is up and running.

Setting Preferences

To set VxVM view and dialog preferences, select **Tools > Preferences >** to display the Preferences window. Then select the **Volume Manager General** tab to show the Preferences. When you have made your changes:

- ◆ Press **Apply** to apply the changes.
- ◆ Press **Reset** to revert back the changes.
- ◆ Press **OK** to apply and dispose the Preference window.
- ◆ Press **Cancel** to discard and close the Preference window.

Note In VxVM releases prior to VxVM 4.0, you could disable the wizards for certain functions by deselecting the Enable Wizard Mode. In this release or VxVM, you can no longer disable wizards.

Show Dialog Top Information Panel

Check this box if you want to show information about an object selected in the tree while you are viewing its related objects in the table.

Launch Task Dialogs as Modal

Check this box if you want to prevent access to other dialogs until the task dialog box is closed.


Show Rules Page

Check the Show rules page box if you want to enable the rules page. By default, the rules page is disabled, and ISP applies the default rules for volume and template creation, and snapshot and transform wizards. (Refer to Set Configuration Rules.)



If you prefer to set the rules yourself, check the box to enable the rules page.

Show Display Name of Capability

Check this box to show the capability display names (shown by default). (Refer to the Select Capabilities for more information on Capability settings.)



Setting Preferences



Disk Tasks

3

VxVM VEA provides you with easy-to-use interfaces for performing tasks and procedures associated with:

- ◆ Managing Disks
- ◆ Managing Disk Groups

The Disk Task Roadmap summarizes these tasks.

Managing Disks

The following tasks are associated with disk management:

- ◆ Adding a Disk to VxVM
- ◆ Scanning Disks
- ◆ Commands for Hot Relocation
- ◆ Reserving a Disk
- ◆ Renaming a Disk
- ◆ Onlining a Disk
- ◆ Mirroring a Disk
- ◆ Evacuating a Disk
- ◆ Disconnecting a Disk
- ◆ Replacing a Disk
- ◆ Recovering Volumes on a Disk
- ◆ Removing a Disk from a Disk Group
- ◆ Forcing Initialization of a Disk
- ◆ Viewing Disk Properties
- ◆ Disabling a Controller

Managing Disk Groups

- ◆ Enabling a Controller
- ◆ Renaming an Enclosure
- ◆ Resizing Disks
- ◆ Setting Disk Usage

Managing Disk Groups

The following tasks are associated with disk group management:

- ◆ Creating a Disk Group
- ◆ Upgrading a Disk Group Version
- ◆ Renaming Disk Groups
- ◆ Deporting Disk Groups
- ◆ Importing Disk Groups
- ◆ Recovering Volumes in Disk Groups
- ◆ Clear Hot Relocation Info
- ◆ Destroying a Disk Group
- ◆ Moving Objects Between Disk Groups
- ◆ Splitting A Disk Group
- ◆ Joining Disk Groups
- ◆ Viewing Disk Group Properties
- ◆ Organizing Disk Groups

Disk Task Roadmap

The following table provides a summary of common VxVM disk and disk group tasks.

To perform this task ...	Use this procedure ...
<i>Disk Setup Tasks:</i>	
Place new disks under VERITAS Volume Manager (VxVM) control	Adding a Disk to VxVM
Set up new disk groups	Creating a Disk Group
<i>Precautionary Tasks:</i>	
Place the boot (root) disk under Volume Manager control	Adding a Disk to VxVM
Make an alternate boot (root) disk	Mirroring a Disk
<i>Maintenance Tasks:</i>	
Place additional disks under VxVM control	Adding a Disk to VxVM
Add more disk space to a disk group	Adding a Disk to VxVM
Create additional disk groups	Creating a Disk Group
Upgrade a disk group	Upgrading a Disk Group Version
<i>Recovery Tasks:</i>	
Move volumes off failing disks	Evacuating a Disk
Replace failed disks	To replace a disk drive
Attempt to recover all volumes on a disk	Recovering Volumes on a Disk
Attempt to recover all volumes in a disk group	Recovering Volumes in Disk Groups
<i>Other Tasks:</i>	
Display the table of disks	Viewing the Table of Disks
Display the table of disk groups	Viewing the Table of Disk Groups
Change a disk name	Renaming a Disk
Change a disk group name	Renaming Disk Groups
Move objects between disk groups	Moving Objects Between Disk Groups
Split a disk group	Splitting A Disk Group
Join disk groups	Joining Disk Groups
Restore access to an offline disk	Onlining a Disk

Adding a Disk to VxVM

To perform this task ...	Use this procedure ...
Temporarily prevent access to a disk group	Deporting Disk Groups
Restore access to a deported disk group	Importing Disk Groups
Remove a disk from VxVM control	Removing a Disk from a Disk Group
Destroy a disk group (permanently)	Destroying a Disk Group
Create an organized disk group	Organizing Disk Groups

Adding a Disk to VxVM

New disks must be set up on the system, placed under VxVM control, then added to a disk group before they can be used for volumes. The **Actions > Rescan** task scans for new devices.

When you place a disk under VxVM control, the disk is either encapsulated or initialized. Encapsulation preserves any existing data on the disk in volumes. Initialization destroys any existing data on the disk. (See the *VERITAS Volume Manager Administrator's Guide* for more information on disk encapsulation and initialization.)

This task assigns a default VxVM disk name to the disk. Once the disk is under VxVM control, you can use the **Rename Disk** task (Renaming a Disk) to change the VxVM disk name (see "Renaming a Disk" on page 45).

Adding a Disk to a Disk Group

Select the uninitialized disk to be placed under VxVM control.

1. Choose **Actions > Add Disk to Disk Group**. The Add Disk to Disk Group wizard appears. Click **Next** to continue.
2. Complete the Add Disk to Disk Group wizard as follows:

Disk group name:	Select the disk group into which the disks will be added. To add the disk(s) to a new disk group, select New disk group , and enter the name of the new disk group in the new dialog box.
Available disks: Selected disks:	Move the disks to be added from Available disks to Selected disks.

When you have provided all necessary information in the dialog box, click **Next**.

3. When the confirmation dialog box appears, click **Yes** to confirm your selection.
4. Click **Finish** to add the disk to the selected group.
5. You now indicate whether the disk should be initialized or encapsulated. (It is necessary to reboot the system before encapsulation takes effect.)

Notes:

- ◆ The **Add Disk to Disk Group** task sets up disks for VxVM use. Disks must be under VxVM control and in a disk group before they can be used to create volumes.
- ◆ Disks must be online before they can be added to a disk group.
- ◆ Disks that already belong to a disk group cannot be added to another disk group.
- ◆ Disks cannot be added to deported disk groups.

Scanning Disks

When VxVM rescans disks, it scans all attached disks for disk configuration changes. It also updates information on all removable media.

Rescanning disks can take several minutes, depending on the number of hardware devices installed.

- ◆ You should rescan disks when you remove disks from or add disks to a computer.
- ◆ From the **Actions** menu, select **Rescan**.

A dialog box comes up with the message “Getting hardware configuration. Please wait.” The box also displays a progress bar showing the percentage of completion for the rescan process.

Reserving a Disk

The **Set Disk Usage** feature can be used to designate a disk as a reserved disk. A reserved disk is not considered part of the free space pool. If you perform a task that requires disk space and you do not specify a disk, VxVM VEA allocates space from disks other than the reserved disk.

▼ To reserve a disk

1. Follow the instructions in Adding a Disk to VxVM to place the disk under VxVM control and in a disk group.
2. Select the disk to be reserved.
3. Choose **Actions > Set Disk Usage**.
4. In the Set Disk Usage window, check the **Reserve Disk** option. Click **OK**.

(To cancel the disk reservation, open the Set Usage window, and uncheck the **Reserve Disk** option.)

Notes:

- ◆ Before you select the **Reserve Disk** option, make sure the disk is under VxVM control and in a disk group.
- ◆ A reserved disk is not available for general use.
- ◆ A reserved disk can only be used if you specify the reserved disk when you perform a task.

Renaming a Disk

The **Rename Disk** procedure changes the VxVM name assigned to a disk. The VxVM disk name is the name given to the disk by the administrator or VxVM. This procedure does not change the physical disk name (device name).

▼ To rename a disk

1. Select the disk to be renamed.
2. Choose **Actions > Rename Disk**.
3. Complete the Rename Disk dialog box as follows:

Disk Name:	If the correct disk name does not already display in this field, enter the disk name to be changed.
New Name:	Enter the new VxVM disk name.

When you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The **Rename Disk** task requires the current disk name and a new disk name.
- ◆ The new disk name must be unique within the disk group.
- ◆ The disk must be in an imported disk group.

Onlining a Disk

The **Online Disk** procedure restores access to a disk that has been taken offline. The disk is made accessible to VxVM again. After bringing a disk back online, the disk must be added to a disk group before it can be used for volumes (see “Adding a Disk to VxVM,”).

▼ To bring a disk online

1. Select the disk to be brought online.
2. Choose **Actions > Online Disk**.
3. Click **Yes** when the confirmation message appears.

Notes:
<ul style="list-style-type: none">◆ When a disk is brought online, the disk is accessible again.◆ Only disks in the offline state can be brought online.

Mirroring a Disk

The **Mirror Disk** procedure mirrors (copies) all concatenated and striped volumes on a disk onto another disk. You can use this procedure to mirror any disk that contains concatenated or striped volumes.

▼ **To mirror all volumes on a disk**

1. Select the disk that contains the volumes to be mirrored onto another disk.
2. Choose **Actions > Mirror Disk**.
3. Enter the target disk name in the Mirror Disk. Click **OK**.

Notes:

- ◆ The **Disk Mirror** task requires a disk name.
- ◆ Only concatenated and striped volumes that reside on a single disk are mirrored onto the other disk. Any RAID-5 or spanned volumes on the disk are ignored.
- ◆ Only disks in the same disk group can be used to create the new mirrors.
- ◆ The target disk must contain sufficient space to accommodate the mirrors.
- ◆ A new mirror cannot be created on a disk that already contains a copy of the volume.
- ◆ If no disks are assigned, VxVM uses available disk space to create the mirrors.

Evacuating a Disk

This feature is similar to disk replacement, but it involves moving the contents of the volumes on one disk to another disk. If a disk begins to fail, you can attempt to protect/preserve the volumes on that disk by evacuating the disk. You can also evacuate a disk if you plan to remove the disk or use the disk elsewhere. The other disk does not have to be empty. The active redundant volumes on the replacement disk will be synchronized automatically.

To Evacuate a disk:

1. Right click on the disk you wish to evacuate.
2. Select **Evacuate Disk**.
3. The Evacuate Disk dialog will appear. Select either “Auto Assign destinations disks” or “Manually assign destination disks”.
 - a. If you chose “Manually assign destination disks” add the destination disks by clicking on the disk and clicking the **Add** button.

Remove disks from the selected destination disks by clicking on the disk and clicking the **Remove** button.
4. Click the **OK** button to evacuate the disk.

When you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The contents of a disk must be evacuated/moved to another disk in the same disk group with sufficient free space.
- ◆ If no target disk is specified, VxVM uses an available disk with sufficient free space.
- ◆ If the disk being evacuated contains part of a mirrored, striped, or RAID-5 volume, do not move the contents of the disk to another disk containing a copy of the mirrored volume or part of the striped/RAID-5 volume.

Disconnecting a Disk

The **Disconnect Disk** procedure disconnects a VxVM disk from its underlying physical disk. You can disconnect a disk in preparation for the replacement of a physical disk. After disconnecting a disk, you can use the Replacing a Disk task to reconnect the disk (see “Replacing a Disk” on page 50).

▼ To disconnect a disk

1. Select the disk to be disconnected.
2. Choose **Actions > Disconnect Disk**.
3. Press **OK** to confirm disconnection.

Notes:

- ◆ Disconnecting a disk that contains volumes can result in loss of data or redundancy.

Replacing a Disk

Disks most often need replacing when they fail or start to behave strangely. A failed or removed disk can be replaced, and active redundant volumes (mirrored and RAID-5) on the replaced disk will be resynchronized automatically after disk replacement.

Failed volumes on the replaced disk, however, will stay failed after disk replacement because no valid data is available for copy. These underlying steps are handled automatically when the disk replacement is performed using VxVM VEA.

Caution This operation alters the partitioning of a disk.

The disk replacement procedure can only be performed on a disk that has failed or is disconnected. VxVM VEA indicates a failed disk by renaming it "Missing Disk". If the disk replacement is successful, the replacement disk will take on the attributes of the failed disk, including the Disk Name.

Note This option is available only on missing or disconnected disks.

To replace a disk drive

1. Right click on the missing disk.
2. Select **Replace Disk**.
A Dialog box appears containing a list of new disks for replacement.
3. Select (click on) the disk you wish to use to replace the missing disk.
4. Click **OK** to perform the Disk Replacement.

Notes:

- ◆ The **Replace Disk** task requires the name of the disk to be replaced and the replacement disk name.
- ◆ The disk to be replaced must be disconnected or removed from its disk group. If the disk has not failed, remove the disk from its disk group before you replace the disk.
- ◆ The new disk can be a disk in the free disk pool or an uninitialized disk. If the new disk is uninitialized, the **Replace Disk** task places the disk under Volume Manager control.
- ◆ The new disk is placed in the disk group of the old disk.

Recovering Volumes on a Disk

The **Recover Disk** procedure performs any necessary volume recovery operations on the volumes on a given disk. The recovery operations depend on the types of volumes on the disk and include starting disabled volumes, resynchronizing mirrors in mirrored volumes, and resynchronizing parity in RAID-5 volumes. After successful recovery, the volumes are again available for use.

Alert icons and the Alert Monitor window may provide clues to let you know that volume recovery is needed.

▼ To recover all volumes on a disk

1. Select the disk containing the volumes to be recovered.
2. Choose **Actions > Recover Disk**.
3. Click **OK** to complete the recover disk procedure.

Note:

In some cases, recovery is not possible. If the volume recovery fails, you can attempt to restore the volume from backup.

Removing a Disk from a Disk Group

You can remove a disk from a disk group if there are no volumes on the disk. If volumes exist, you must delete or evacuate them before you can revert the disk. Disks shown in Volume Manager to be **Unreadable** (inaccessible due to corruption) as well as **Foreign** (a disk that has moved to another machine) can be removed, but doing so causes you to lose all existing data on the disk that is removed from the group. If a disk has failed and is marked as **Missing**, you can also remove it from the disk group.

To remove a disk from a disk group, use the **Remove Disk from Disk Group** command.

The steps for removing a disk from a disk group are as follows:

1. Right-click on the selected disk or on the Disk Group and select **Remove Disk from Disk Group**.
2. Make sure the proper disk or disks that you want to remove are showing in the right pane of the window and click **Next** to continue.

The **Remove Disk from Disk Group** procedure removes the disk from VxVM control. If the disk to be removed contains volumes, you must evacuate the disk before you remove the disk.

When you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ If the disk contains a volume that has only one mirror and either has a mounted file system or is used as a swap area, you must evacuate the disk before you remove it.
- ◆ When a disk is removed from a disk group or VxVM control, the disk cannot be accessed.
- ◆ A disk that is in use (that is., contains volumes) must not be removed. Removing a disk with volumes can result in loss of data or data redundancy. If volumes exist on the disk to be removed, use the **Evacuate Disks** task to move the volumes to another disk.
- ◆ Any non-redundant volumes on the removed disk become disabled.

Forcing Initialization of a Disk

Caution Forcing the initialization or reinitialization of a disk can result in data loss.

The **Initialize Disk** procedure forces the initialization of a disk. This task is intended for use in troubleshooting disks that have already been placed under VxVM control. To initialize a disk for VxVM use, use the **Add Disks** task instead (see Adding a Disk to VxVM).

▼ **To (re)initialize a disk**

1. Select the disk to be initialized.
2. Choose **Actions > Initialize Disk**.
3. Complete the **Initialize Disk** dialog box as follows:

Disk Name:	If the correct disk name does not already display in this field, enter the disk name.
------------	---------------------------------------------------------------------------------------

When you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The **Initialize Disk** task forces the (re)initialization of a disk and can result in data loss.
- ◆ The **Initialize Disk** task overwrites existing disk header and configuration information on the disk. Any volume copies or regions located on the disk are not accessible after you reinitialize the disk.

Resizing Disks

You can resize (shrink, or grow) a virtual disk device with a SCSI interface presented by a smart switch, smart array or RAID controller. The device must be online in an imported disk group. The resize function can be performed only on a host licensed for the Dynamic LUN Expansion option.

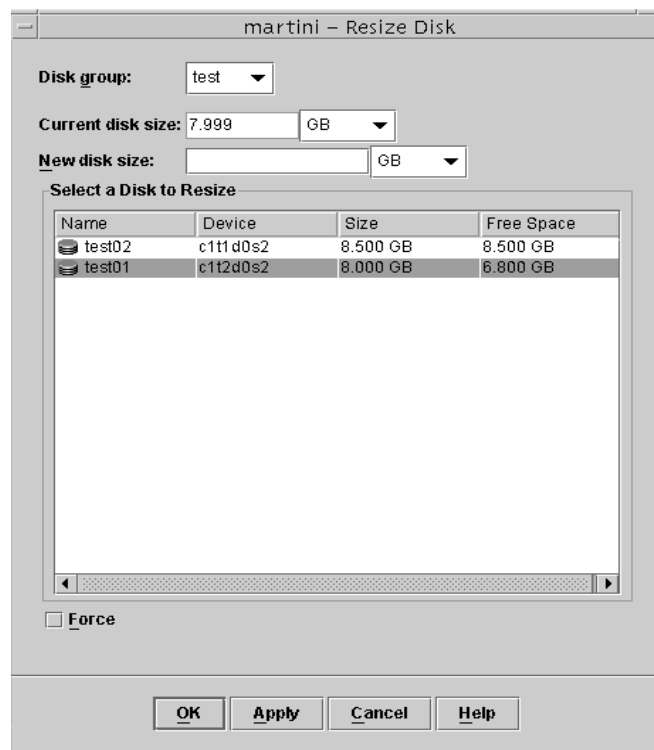
The data is preserved. Only sectors at the end of device can be added or removed.

Note You must ensure that this operation is only invoked for storage devices that preserve the data while resizing.

To resize a disk:

1. Select the disk you wish to resize.
2. Choose **Actions -> Resize Disk....** The Resize Disk dialog appears.

Resize Disk Dialog



3. Complete the Resize Disk Dialog Box as follows:

Disk group:	You can select a disk group other than the one originally chosen.
Current disk size:	The current size of the selected disk.
New disk size:	If the new size is smaller than the current disk size, then this is a shrink request and the subdisks fall outside the new disk size need to be preserved.
Select a Disk to Resize:	Select the disk to be resized (from the selected disk group)
Force	Before reducing the size of a device, any volumes on the device should first be reduced in size or moved off the device. By default, the resize fails if any subdisks would be disabled. Selecting the Force checkbox overrides this behavior, but this can result in loss of data.

4. After you have provided all necessary information in the dialog box, click **OK**.

Setting Disk Usage

Use the Set Disk Usage feature to exclude a disk from certain operations by marking it as a spare, excluding it from Hot Relocation, or marking it as a reserved disk, or as a disk reserved for ISP use only.

To set disk usage:

1. Select the disk you wish to modify.
2. Choose **Actions -> Set Disk Usage....**
3. The Set Disk Usage dialog appears. Select the options you wish to implement.
 - Spare: (to be used as a spare disk in case it's needed)
 - No Hot use: (cannot be used as a Hot Spare or for Hot Relocation)
 - Reserved: (reserved for other needs)
 - Reserved for Allocator: (reserved for ISP use only)
4. After you have provided all necessary information in the dialog box, click **OK**.

Viewing the Table of Disks

You can display the table of disks under VxVM control by clicking on the Disk node. The following information is displayed.

Field	Description
Device	Lists the disks under VxVM control, by device name, for this host. You can also use this field to: <ul style="list-style-type: none">◆ Change the fields displayed.◆ Sort devices in ascending or descending order.
Internal Name	The name you have defined for the disk, or the default name allocated to the disk by VxVM.
Group Name	The name of the group to which the disk is added.
Pool Name	The name of the storage pool in which a system allocator disk resides.
Status	Disk status can be Imported, Deported, Not Initialized.
Size	Size of disk.
Type	Disk type.
% Used	The percentage of the disk space already used. This is color-coded, according to whether it shows Normal Usage, High Usage, or Critical Usage.
Spare	Indicates whether or not the disk has been designated as a spare for use with hot relocation.
Unallocated Space	Amount of free space available to create a partition or volume.
Reserved	Indicates whether the disk is reserved Not to be used as a disk.
Hot Use	Indicates whether the disk is not to be used as a hot spare or to receive volumes from the Hot Relocation procedure.
UDID	Indicates the Unique Disk Identifier (UDID) of the disk.
WWN	.Indicates the World Wide Name of the disk.

Viewing Disk Properties

The Disk Properties window contains detailed information about the selected disk. The following information is displayed.

Category	Description
Internal Name	Name of disk in VxVM (default name is <i>groupnamexx</i>).
Pool Name	Name of the Storage Pool (data pool or clone pool).
CDS	Yes: CDS disk; No: Not CDS.
Status	Fields values can be Imported, Deported, Not Initialized.
Capacity	Size.
Unallocated Space	Amount of free space that is available to create a partition or volume.
UDID	Unique disk identifier. This is a DDL-related property. DDL is the Device Discovery Layer that handles all disk discovery related functions needed for the VxVM/DMP operation.
WWN	This is the World Wide Name of the disk. This is a DDI-related property. DDI properties are device driver properties primarily used in Solaris OS.
Spare	Indicates whether (Yes) or not (No) the disk has been designated as a hot spare.
Reserved	Indicates whether the disk is reserved Not to be used as a disk.
Hot Use	Indicates whether the disk is not to be used as a hot spare or to receive volumes from the Hot Relocation procedure.
Allocator Use	Indicates an ISP disk.
Comment	You can enter up to 40 characters of comment in this area.
Spare	Indicates whether (Yes) or not (No) the disk has been designated as a hot spare.

Disabling a Controller

The controller **Disable** procedure disables I/O to a controller.

Caution If your system supports the Dynamic Multipathing (DMP) feature; it should be noted that performing the controller **Disable** procedure will prevent DMP from accessing the paths associated with a controller.

You can disable access to a controller in preparation for hardware maintenance.

▼ **To disable a controller**

1. Select the controller to be disabled.
2. Choose **Actions > Disable**.
3. When you have checked the information in the dialog box, click **OK**.

Enabling a Controller

The controller **Enable** procedure enables I/O to a controller. (If your system supports the DMP feature, and if DMP is enabled, then the **Enable** procedure also restores access to the paths associated with a controller.)

▼ **To enable a controller**

1. Select the controller to be enabled.
2. Choose **Actions > Enable**.
3. When you have checked the information in the dialog box, click **OK**.

Renaming an Enclosure

The **Enclosure Rename** procedure changes the name of an enclosure (disk array).

▼ **To rename an enclosure**

1. Select the enclosure to be renamed.
2. Choose **Actions > Enclosure Rename**.
3. Complete the Rename Enclosure dialog box as follows:

Enclosure Name:	If the correct enclosure name does not already display in this field, enter the name.
New Name:	Enter the new enclosure name.

When you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The **Enclosure Rename** task requires the current enclosure name and a new enclosure name.
- ◆ The new enclosure name must be unique within the disk group.

Creating a Disk Group

Disks must be placed in disk groups before they can be used by VxVM. You can create disk groups to organize your disks into logical sets of disks.

When you place a disk under VxVM control, the disk is initialized. Initialization destroys any existing data on the disk.

▼ To create a disk group

1. Select the Disk Groups folder.
2. Select **Actions > New Disk Group**. The New Disk Group Wizard welcome screen appears.
3. You can suppress the Welcome to Create New Disk Group page by selecting the **Do not show this page next time** checkbox.
4. Click **Next** to continue.
5. The New Disk Group Wizard appears. Provide information as described in the following table.

Group name:	Enter a new disk group name in the Group Name field.
Create cluster group	If the cluster feature is available, you can select the Create cluster group checkbox if the new disk group is to be used with clusters.
Available disks: Selected disks:	Select which disks (from Available disks:) you want to include in the group. Make sure the disks you want to include are in the right pane of the window (Selected disks:). (Normally, you would add all the disks you want in the group at this point. However, you can always add more disks later with the Add Disk to Disk Group procedure.)
Disk names:	Type in the names of any disks to be added that do not appear under Selected disks:

6. When you have provided all necessary information in the dialog box, click **Next>**.

7. If the disks are not empty, you now need to indicate whether the disk should be initialized or encapsulated. The initialize option is not available for disks with mounted partitions or swap areas. If you initialize the disk, any existing data on the disk will be destroyed. If you encapsulate the disk, any existing data will be preserved in volumes.
8. If you chose to encapsulate the disk, the disk needs to be rebooted before encapsulation can take effect. You have the option of rebooting the system now, or later.
9. Click **Next** to add the disks to the New Disk Group.
10. The Organization principle screen now appears.
11. Click **None**, if you do not want to organize the disk group. Otherwise, click any of the other organization categories under **Organize DiskGroup By** to create an ISP disk group (refer to Organizing Disk Groups).
12. If you choose an organization principle other than **None**, the Specify Pool Names screen appears. Specify the names for the **Data pool** and the **Clone pool** and click **Next**. You must specify the pool names in order to create the first storage pool for an ISP disk group.
13. The next screen confirms the disks you have selected. Click **Finish** to continue if you are satisfied with the disk selection. If you are not satisfied, you can click the **<Back** button to go back to the previous screen in order to modify your selections.
14. After clicking **Finish**, the new disk group will then appear under the Disks Groups node.

Notes:

- ◆ The disk group name must be unique.
- ◆ The new disk group must contain at least one disk.
- ◆ Only disks that are online and do not belong to a disk group can be used to create a disk group.

Upgrading a Disk Group Version

The **Disk Group Upgrade** procedure upgrades a disk group to the disk group version that is compatible with the current VxVM release. Some new features and tasks only work on disk groups with the current disk group version, so you must upgrade existing disk groups before you perform these tasks. However, after you upgrade a disk group, you cannot access that disk group from machines running earlier releases of VxVM that do not support a higher current disk group version.

To determine if a disk group needs to be upgraded, select the disk group and open the Disk Group Properties window. The **Current Version** field of the Properties window indicates whether the disk group has been upgraded to the current disk group version (**No** indicates that disk group needs to be upgraded).

▼ To upgrade a disk group

1. Select the disk group to be upgraded to the current version.
2. Choose **Actions > Upgrade Disk Group Version**.
3. Confirm the upgrade by clicking **OK**

Notes:

- ◆ If you have disk groups created with earlier releases of VxVM, these must be upgraded before they can take advantage of new features and tasks in the current release.
- ◆ If a disk group is upgraded to the current disk group version, you cannot import that disk group on a machine running an earlier VxVM release. Disk groups that must remain backward compatible with earlier releases of VxVM must not be upgraded.

Renaming Disk Groups

The **Rename Disk Group** procedure changes the name of an existing disk group. You can rename a disk group to provide a more appropriate name or to avoid a name conflict with a disk group to be imported.

▼ To rename a disk group

1. Select the disk group to be renamed.
2. Choose **Actions > Rename Disk Group**.
3. Complete the Rename Disk Group dialog box.
4. When you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The **Rename Disk Group** task requires the current disk group name and a new disk group name.
- ◆ The new disk group name must be unique.
- ◆ The **Rename Disk Group** task deports and re-imports the disk group. If volumes in the disk group are in use, the deport fails and the disk group is not renamed.
- ◆ The **Rename Disk Group** task updates the file system table file, if necessary.
- ◆ If the disk group contains volumes with mounted file systems, you may need to unmount the file systems.

Deporting Disk Groups

The **Deport Disk Group** procedure makes a disk group and its volumes inaccessible. Once deported, a disk group cannot be used until it is imported. You can deport a disk group to make the disk group *temporarily* unavailable.

While a disk group is deported, you cannot remove and reuse any of its disks. To remove disks, you must import the disk group and then either destroy the disk group (to remove all of its disks) or remove selected disks from the disk group.

▼ To deport a disk group

1. Select the disk group to be deported.
2. Choose **Actions > Deport Disk Group**.
3. If you want to change the disk name and define a new host, then check **Deport options** and complete the **New name** and **New host** fields - type in the name of the new name of the disk group and the name of the host you are deporting the disk group to.
4. Click **OK** to complete the deportation.

Notes:

- ◆ A deported disk group cannot be accessed. To access a deported disk group, the disk group must be imported.
- ◆ A disk group cannot be deported if any volumes in that disk group are in use (open).
- ◆ When a disk group is deported, the host ID (the name of the machine where the disk group currently resides) stored on all disks in the disk group is cleared and the disk group is not re-imported automatically when the system is rebooted. However, if you specify a host in New host option, the specified host imports the disk group at reboot.

Importing Disk Groups

The **Import Disk Group** procedure makes a deported (inaccessible) disk group and its volumes accessible again. To import a deported disk group, you must know the former name of the disk group, and this name must not be in use. In addition, at least one disk name formerly assigned to the deported disk group must not be in use.

▼ To import a disk group

1. Select the appropriate host machine node.
2. Select the disk group for which the group is to be imported.
3. Choose **Actions > Import Disk Group**.
4. Complete the Import Disk Group dialog box as follows:

New Name:	Enter a new disk group name in the New Name field. to change the name of the disk group at import.
Clear host ID	Select Clear host ID to clear the existing host ID stamp (name of the host machine that currently "owns" the disk group) on all disks in the disk group at import. Do not use this option if another host is using any disks in the disk group.
Force	Select Force to force the disk group import when the host cannot access all disks in the disk group This option can be used to import a disk group that contains a failed disk, but can lead to disk group inconsistency if all disks are still usable.
Start all volumes	Select Start all volumes checkbox to start all volumes associated with the disk group on import.

Recovering Volumes in Disk Groups

When you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ Only deported disk groups can be imported.
- ◆ A deported disk group cannot be imported if another disk group with the same name has been created since the disk group was deported.
- ◆ When a disk group is imported, the system stamps its host ID (machine name) on all disks in the disk group. A disk group import fails if one of the disks is stamped with a host ID that does not match the others. This ensures that dual-ported disks cannot be managed (and possibly corrupted) by two systems at the same time. If you are *sure* that the disk group is *not* in use by another host, you can clear the host IDs and import the disk group by selecting the **Clear host ID** option.

Recovering Volumes in Disk Groups

The **Recover Disk Group** procedure performs any necessary volume recovery operations on the volumes in a given disk group. The recovery operations depend on the types of volumes in the group and include starting disabled volumes, resynchronizing mirrors in mirrored volumes, and resynchronizing parity in RAID-5 volumes. After successful recovery, the volumes are again available for use.

Alert icons and the Alert Monitor window may provide clues to let you know that volume recovery is needed.

▼ **To recover all volumes in a disk group**

1. Select the disk group containing the volumes to be recovered.
2. Enter the disk group name if the correct disk group name is not shown in the dialog box.
3. Choose **Actions > Recover Disk Group**.
4. Click **OK**.

Note:

In some cases, recovery is not possible. If the volume recovery fails, you can attempt to restore the volume from backup.

Commands for Hot Relocation

The Hot Spare and Hot Relocation procedures enable a system to automatically react to I/O failures on redundant (mirrored or RAID-5) volumes, and to restore redundancy and access to those volumes. The following commands can be used to perform relocation tasks:

- ◆ Undo Hot Relocation
- ◆ Clear Hot Relocation Info
- ◆ Customize Disks for Hot Relocation

Note The Hot Relocation and Hot Spare operations are 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.

About Hot Spare and Hot Relocation

One or more disks in each disk group can be designated as hot spares by using the “set disk usage” command. These hot spare disks provide a location where data can be moved from failed disks or volumes.

In Hot Spare mode, all subdisks containing redundant volumes will be moved together from a failed disk to a spare disk. In order for the Hot Spare operation to occur, there must be enough free space on the spare disk for all the redundant volumes from the original disk.

When a disk fails, hot relocation automatically moves all subdisks containing redundant volumes from the failed disk to hot spare disks. If no disks have been designated as spares or there is not enough spare disk space, a combination of spare space and free space is used. Depending on available disk space, the subdisks are not always all moved to the same disk, but may be scattered within the disk group.

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.

When selecting space for relocation, hot relocation preserves the redundancy characteristics of the Volume Manager 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 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.

Undo Hot Relocation

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

Undo Hot Relocation

The **Undo Hot Relocation** procedure relocates subdisks back to their repaired original disk or replacement disk and restores a system to its original configuration. If hot relocation scattered subdisks from a failed disk to several disks within a disk group, the undo hot relocation command will summon all of them back to a single disk without you having to find and move each one individually.

Note The **Undo Hot Relocation** procedure is only available after a hot relocation or hot sparing procedure has occurred.

1. Right-click on the disk group that contains the hot relocated volumes.
2. Select **Actions > Undo Hot Relocation**.
A window appears showing the disks that failed and from which subdisks were hot relocated.
3. Select the disks you wish to unrelocate by clicking on the corresponding circle in the column under the check mark.
4. Click **OK**. The hot relocated subdisks will return to the original disk.

Clear Hot Relocation Info

The **Clear Hot Relocation Info** command for disk groups deletes the information that describes the original location of any subdisks that have been hot relocated. Use this command if you do not plan to move the subdisks back to their original configuration.

1. Right-click on a disk group that contains hot relocated subdisks.
2. Select **Actions > Clear Hot Relocation Info**.
3. A dialog box will appear asking if you are sure you want to clear the relocation information on the disk group. Click **Yes** if you want to clear the hot relocation information.

Caution Once you clear the Hot Relocation Information, you will not be able to perform an Undo Hot Relocation on that disk group.

Notes:

- ◆ The **Clear the Relocation Info** procedure requires the name of the original disk (even if an alternate disk is specified).
- ◆ When a disk fails, the hot-relocation feature attempts to move subdisks from the failed disk to other disks. The **Undo Relocation Info** procedure can be used to return the relocated subdisks to their original disk locations after the original disk is repaired or replaced.
- ◆ It is not possible to return relocated subdisks to their original disks after their disk group relocation information has been cleared with the **Clear the Relocation Info** task.

Customize Disks for Hot Relocation

You can customize your disks in the following ways:

- ◆ Specify individual disks to act as Hot Spares
- ◆ Prevent a disk from being used for Hot Relocation or Hot Spare
- ◆ Mark a disk as Reserved
- ◆ Mark a disk as Reserved for Allocator

Each disk must be customized individually, you cannot set disk usage for a group of disks:

Customize Disks for Hot Relocation

1. Click on the disk to be customized.
2. Choose **Actions > Set Disk Usage**.

A window appears offering the customization options. By default, no options are selected.

The reserved disk option (**Reserved** or **Reserved for Allocator**) allows a disk to be marked so that it will not be used in automatic selection operations. For example, if a disk is marked as reserved during the create new volume wizard using the default option to auto select disks, the volume would not be created on the reserved disk. However, if you wish to create a volume on that specific disk, you could choose the option to manually select the disks for the new volume and then select that disk.

3. A disk can be marked for use for spare only (**Spare**), or for neither hot spare or hot relocation (**No hot use**). When you are satisfied with your selection, click **OK**.

You can review the settings for all your disks by selecting the Disks folder in the tree pane and then viewing Disks tab in the right pane. The settings for each disk are listed in the **Spare**, **Reserved**, **Reserved for Allocator**, and **Hot Use** columns.

Destroying a Disk Group

Caution Destroying a disk group can result in data loss.

The **Destroy Disk Group** procedure *permanently* removes a disk group from VxVM control. This procedure also reinitializes all of the disks in the disk group as empty disks and places them in the free disk pool for reuse.

Destroy a disk group *only* if you are sure that you no longer need the volumes and data in the disk group.

▼ To destroy a disk group

1. Right-click on the disk group to be destroyed.
2. Choose **Destroy Disk Group**.
3. When the warning prompt appears; select **Yes** to confirm that you want all volumes and disks removed from the disk group.
4. The disk group is then destroyed.

Notes:

- ◆ Use the **Destroy Disk Group** task with caution. Destroying a disk group permanently removes the disk group from VxVM control and can result in data loss.
- ◆ When a disk group is destroyed, its disks are returned to the free disk pool. Any volumes in the disk group are removed.
- ◆ A disk group cannot be destroyed if any volumes in that disk group are in use (open) and/or contain mounted file systems.

Moving Objects Between Disk Groups

The following procedure moves a self-contained set of VxVM VEA objects between imported disk groups. This operation is performed on VxVM VEA objects such as disks or top-level volumes, and includes all component objects such as sub-volumes, plexes and subdisks.

▼ **To perform a disk group move operation:**

1. Right-click on the disk group to be moved.
2. Select Move **Disk Group**.
3. Complete the Move Disk Group Dialog Box as follows:

Disk Group Name:	Select the name of the disk group from which objects are to be moved from the Disk Group Name pulldown menu.
Target Disk Group Name:	Enter the name of the disk group into which objects are to be moved.
Select Objects to move:	<p>Enter the objects to be moved.</p> <ul style="list-style-type: none"> ◆ To select volumes to move, click Move disk group by volumes. To select all of the listed volumes, click Add All. ◆ To select disks to move, click Move disk group by disks. To select all of the listed disks, click Add All. ◆ To select storage pools to move, click Move disk group by storage pool. To select all of the listed storage pools, click Add All. <p>Click on OK.</p>
Expand:	Click Expand to specify that the objects to be moved should include all other disks that contain subdisks that are associated with the specified objects or with objects that they contain. For example, if you specify that disk01 should be moved and disk01 contains vol1 which is also located on disk02, then by using the Expand button, you ensure that disk02 will also be moved.
Show Dependencies:	Click Show Dependencies to display the objects affected by the move.

Notes:

- ◆ Before moving volumes between disk groups, stop all applications that are accessing the volumes, and unmount all file systems that are configured in the volumes.
- ◆ The objects to be moved must be self-contained, meaning that the disks that are moved must not contain any other objects that are not intended for the move.
- ◆ This operation fails if it would remove all the disks from the source disk group.
- ◆ Volume states are preserved across the move.
- ◆ The reconfiguration must involve an integral number of physical disks.
- ◆ Objects to be moved must not contain open volumes.
- ◆ Data change objects (DCOs) and snap objects that have been dissociated by Persistent FastResync cannot be moved between disk groups.
- ◆ If your system supports VERITAS Volume Replicator (VVR), then it should be noted that VVR objects cannot be moved between disk groups.
- ◆ For a disk group move to succeed, the source disk group must contain at least one disk that can store copies of the configuration database after the move.
- ◆ For a disk group move to succeed, the configuration database in the target disk group must be able to accommodate information about all the objects in the enlarged disk group.
- ◆ The operation can only be performed on the master node of a cluster if either the source disk group or the target disk group is shared.

Splitting A Disk Group

The disk group split operation involves splitting the contents of an imported disk group and creating a new disk group (the target disk group) with the objects defined in the dialog box. This operation is performed on VxVM VEA objects such as disks or top-level volumes, and includes all component objects such as sub-volumes, plexes and subdisks.

▼ **To perform a disk group split operation:**

1. Select the disk group to be split.
2. Select **Actions > Split Disk Group**.
3. Complete the Split Disk Group Dialog Box as follows:

Disk Group Name:	Enter the name of the disk group from which objects are to be split.
Target Disk Group Name:	Enter the name of the target disk group which will consist of the objects to be split.
Select Objects to move:	<ul style="list-style-type: none"> ◆ To select disks to move, click the checkbox Split disk group by disk, and choose from the displayed list. ◆ To select volumes to move, click the checkbox Split disk group by volume, and choose from the displayed list. ◆ To select volumes sets to move, click the checkbox Split disk group by volume sets, and choose from the displayed list. ◆ To select storage pools to move, click the checkbox next to Split disk group by storage pool, and choose from the display list.
Expand:	Click Expand to specify that the objects to be split should include all other disks that contain subdisks that are associated with the specified objects or with objects that they contain. For example, if you specify that disk01 should be split off and disk01 contains vol1 which is also located on disk02, then by using the Expand button, you ensure that disk02 will also be split off.
Show Dependencies	Select the Show Dependencies checkbox to show any other disks that would be dragged to the new disk group along with the selected disks.

Notes:

- ◆ Before carrying out the split operation, stop all applications that are accessing the volumes, and unmount all file systems that are configured in the volumes.
- ◆ The objects to be split off must be self-contained, meaning that the disks that are split must not contain any other objects that are not intended for the move.
- ◆ This operation fails if it would remove all the disks from the source disk group.
- ◆ Volume states are preserved across the split.
- ◆ An existing deported disk group is destroyed if it has the same name as the target disk group.
- ◆ The reconfiguration must involve an integral number of physical disks.
- ◆ Objects to be split off must not contain open volumes.
- ◆ Data change objects (DCOs) and snap objects that have been dissociated by Persistent FastResync cannot be moved between disk groups.
- ◆ If your system supports VVR, then it should be noted that VVR objects cannot be moved between disk groups.

Joining Disk Groups

This procedure joins two disk groups by moving all VxVM objects from an imported source disk group to an imported target disk group. The source disk group is removed when the join is complete. This operation is performed on VxVM VEA objects such as disks or top-level volumes, and includes all component objects such as sub-volumes, plexes and subdisks.

▼ **To perform a disk group join operation:**

1. Select the disk group to be joined.
2. Select **Actions > Join Disk Group**.
3. Use the source disk group name shown or replace it with one from the pull-down menu.
4. Use the target disk group name shown, or replace it with one from the pull-down menu.
5. Click **OK**.

Note:

- ◆ Before performing a join operation, stop all applications that are accessing the volumes, and unmount all file systems that are configured in the volumes.
- ◆ The reconfiguration must involve an integral number of physical disks.
- ◆ Objects to be joined must not contain open volumes.
- ◆ Data change objects (DCOs) and snap objects that have been dissociated by Persistent FastResync cannot be moved between disk groups.
- ◆ If your system supports VVR, then it should be noted that VVR objects cannot be moved between disk groups.
- ◆ For a disk group join to succeed, both the source and target disk groups must contain at least one disk that can store copies of the configuration database after the split.

Viewing the Table of Disk Groups

You can display the table of disks groups under VxVM control by clicking on the Disk Group node.

The following information is displayed.

Field	Description
Name	Lists the disk groups under VxVM control, by name, for this host. You can also use this field to: <ul style="list-style-type: none">◆ Change the fields displayed.◆ Sort disk groups in ascending or descending order.
Status	Group status can be Imported, Deported, Not Initialized.
Size	Size of disk group.
Unallocated	Amount of free space available to create a disk group.
CDS	Cross platform data sharing. Sharing data between heterogeneous systems (e.g., SUN and HP) where each system has direct access to the physical devices used to hold the data.
Type	Disk Group type.
# Disks	Number of disks in the group.
# Volumes	Number of volumes in the group.
Current Version	Indicates whether the disk group has been upgraded to the current disk group version. (No indicates that disk group needs to be upgraded.)
Version	Indicates the disk group version number.

Viewing Disk Group Properties

The Disk Group Properties window contains detailed information about the selected disk group.

- ◆ To view the properties of a disk group, select the disk group and then choose **Properties** from the popup menu.

Organizing Disk Groups

Use the **Organize Disk Group** feature to create a disk group that can have only ISP-style volumes.

▼ To organize a disk group

1. Select the disk group to apply the organization principle.
2. Choose **Actions > Organize Diskgroup**.
3. If the **Welcome screen** appears, click **Next**. You can suppress the Welcome Screen by checking the **Do Not Show This Page Again** checkbox.
4. The organize principle selection screen appears. You can select a different disk group, at this point, if you wish.
5. Select the **Organize Principle** option to be applied to the disk group. Each option has a descriptor associated with it. To view the descriptor for each option, simply click on the option you wish to see the descriptor for and a definition will expand below the descriptor. Once you choose a specific option, you are confined to the specified capabilities within that option.
6. Click **Next** to go to the **Name Storage Pools** page. From the pulldown menu, select the names for your data pool and clone pool.
7. Click **Next** to go to the **Summary Page**. This displays the organize principle selected, the disk group to which the principle is applied, and a description of the disk group.
8. Click **Finish** to apply the organization principle to the selected disk group.

Volume Tasks

4

VxVM Volumes tasks include:

- ◆ Creating Volumes
 - ◆ Creating a VxVM Volume
 - ◆ Creating a VxVM ISP Volume
- ◆ Resizing a Volume
- ◆ Renaming a Volume
- ◆ Adding and Removing Columns in a VxVM ISP Volume
- ◆ Changing a VxVM Volume Layout
- ◆ Changing the VxVM ISP Column Count
- ◆ Adding a Mirror to a Volume
- ◆ Adding a Log to a Volume
- ◆ Stopping and Recovering a Volume
- ◆ Preparing to Restore a Volume From Backup
- ◆ Recovering a Volume
- ◆ Removing a Mirror From a Volume
- ◆ Removing a Log From a Volume
- ◆ Deleting a Volume
- ◆ Attaching Mirrors
- ◆ Moving, Splitting, and Joining Subdisks
- ◆ Viewing Volume Properties

Volume Task Roadmap

The following table provides a summary of common VxVM volume tasks.

To perform this task ...	Use this procedure ...
<i>Volume Setup Tasks:</i>	
Create new volumes	Creating Volumes
<i>Precautionary Tasks:</i>	
Mirror (copy) volumes onto other disks for redundancy	Adding a Mirror to a Volume
Create mirrored volumes	Mirrored
Add a log to a volume	Adding a Log to a Volume
<i>Maintenance Tasks:</i>	
Changing the space associated with a volume	Resizing a Volume
Resizing volumes	Resizing a Volume
Change volume layouts	Changing a VxVM Volume Layout
Change an <i>ISP</i> volume layouts	Changing the VxVM ISP Column Count
Add mirrors to volumes that should be redundant	Adding a Mirror to a Volume
Move portions of volumes to other disk space	Moving a Subdisk
<i>Recovery Tasks:</i>	
Attempt to recover a volume	Recovering a Volume
Attempt to recover all volumes on a disk	Recovering Volumes in Disk Groups (see Chapter 3)
Attempt to recover all volumes in a disk group	Recovering Volumes in Disk Groups (see Chapter 3)
Prepare to restore a volume from backup	Preparing to Restore a Volume From Backup

To perform this task ...	Use this procedure ...
<i>Other Tasks:</i>	
Change a volume name	Renaming a Volume
Restore access to a stopped volume	Recovering a Volume
Remove a mirror from a volume	Removing a Mirror From a Volume
Remove a DRL log from a volume	Removing a Log From a Volume
Remove a RAID-5 log from a volume	Removing a Log From a Volume
Remove a volume	Deleting a Volume
Starting a volume	Recovering a Volume
Temporarily prevent access to a volume	Stopping and Recovering a Volume
Temporarily prevent access to a volume mirror	Removing a Mirror From a Volume

VxVM Volumes

VxVM uses logical volumes to organize and manage disk space. A volume is made up of portions of one or more physical disks, so it does not have the physical limitations of a physical disk.

A volume can provide greater capacity and better availability and performance than a single physical disk. A volume can be extended (grown) across multiple disks to increase capacity, mirrored (copied) on another disk to provide data redundancy, and/or striped across multiple disks to improve I/O performance.

You can use VEA to create the following types of volumes:

- ◆ Concatenated Volume

A concatenated volume is made up of one or more disk regions that are linked together (concatenated) in a linear fashion. A concatenated volume can consist of disk regions on multiple disks; a concatenated volume that extends across two or more disks is also known as a *spanned volume*.

Note Data in this type of volume cannot be recovered if the underlying disk fails. However, a concatenated volume can be mirrored (copied) onto other disks to protect its data against disk failure.

- ◆ Striped Volume

Striped volume data is interleaved (striped) across two or more physical disks. Striped volume data is spread across the disks alternately and evenly in small, equal-sized portions of data called *stripe units*. Striping improves performance.

Note Data in this type of volume cannot be recovered if one of the underlying disks fails. However, a striped volume can be mirrored (copied) onto other disks to protect its data against disk failure.

- ◆ RAID-5 Volume

RAID-5 volume data is interleaved (striped) across three or more physical disks. Within each stripe across the set of disks, the data on one of the disks is parity data. If one of the physical disks fails, the parity data can be used to reconstruct and recover the lost data.

Note RAID-5 volumes cannot be mirrored.

- ◆ **Mirrored Volume**

Volumes with concatenated or striped layouts can be mirrored to increase data availability. All of the data in a mirrored volume is duplicated on at least one other physical disk. If one of the disks fails, the data can still be accessed from one of the remaining disks.

The plexes in a mirrored volume typically have the same layout, but a volume can consist of plexes with different layouts. A mirrored volume has a “mixed” layout if the plexes in the volume have different layouts.

- ◆ **Layered Volume**

A layered volume is built on one or more other volumes. The underlying volumes are typically mirrored. In layered volumes, mirroring is done at a lower level and with smaller granularity than with non-layered volumes, so each mirror covers a relatively small storage region.

Layered volumes tolerate disk failure better than non-layered volumes and provide improved data redundancy. If a disk in a layered volume fails, only a portion of the redundancy is lost and recovery time is usually quicker than it would be for a non-layered volume. Layered volumes also reduce the chance that two disk failures will result in lost data.

The underlying volumes in a layered volume are used exclusively by VxVM and are not intended for user manipulation.

With VEA, you can create the following types of layered volumes:

- ◆ **Concatenated Mirrored Volume**

A concatenated mirrored volume is a layered volume that concatenates several underlying mirror volumes.

- ◆ **Striped Mirrored Volume**

A striped mirrored volume is a layered volume that stripes several underlying mirror volumes.

VxVM Volumes

The following table summarizes the advantages and disadvantages of each volume layout.

Layout Comparison

Layout	Advantages	Disadvantages
Concatenated	<ul style="list-style-type: none"> ◆ The volume can consist of disk regions that are not adjacent. ◆ The volume can span multiple disks. A spanned volume can have a greater capacity than a single disk. ◆ The volume can be mirrored to protect data against disk failure and reduce the risk of data loss. 	<ul style="list-style-type: none"> ◆ There is a risk of data loss if the volume is not mirrored. ◆ Spanning a volume across multiple disks increases the chance that a disk failure will result in the failure of the volume. ◆ Concatenation does not improve I/O performance.
Striped	<ul style="list-style-type: none"> ◆ Striping provides improved read and write performance. ◆ Striping can help to balance the I/O load from multi-user applications across multiple disks. ◆ The volume can be mirrored to protect data against disk failure. 	<ul style="list-style-type: none"> ◆ There is a risk of data loss if the volume is not mirrored. ◆ Striping a volume across multiple disks increases the chance that a disk failure will result in the failure of the volume. ◆ A striped volume requires at least two disks.
RAID-5	<ul style="list-style-type: none"> ◆ RAID-5 provides protection against disk failure and reduces the risk of data loss. ◆ RAID-5 provides data redundancy by storing parity (a calculated value) on the disks. If a disk fails, the parity is used to reconstruct the missing data. ◆ RAID-5 requires less storage space than mirroring. ◆ RAID-5 provides improved read performance. 	<ul style="list-style-type: none"> ◆ RAID-5 provides relatively slow write performance. ◆ RAID-5 volumes cannot be mirrored. ◆ A RAID-5 volume requires at least three disks. If logging is enabled, a RAID-5 volume requires at least four disks. ◆ A RAID-5 volume cannot survive multiple disk failures. However, optional RAID-5 logs reduce this risk. ◆ If one of the volume disks is inaccessible, performance is degraded.

Layout Comparison

Layout	Advantages	Disadvantages
Mirrored	<ul style="list-style-type: none"> ◆ Mirroring provides protection against disk failure and reduces the risk of data loss. ◆ Mirroring provides data redundancy by maintaining multiple complete copies of volume data on different disks. ◆ Mirroring provides improved read performance. 	<ul style="list-style-type: none"> ◆ A mirrored volume requires extra disk space. Each mirror requires enough disk space to contain a complete copy of the volume.
Concatenated Mirrored	<ul style="list-style-type: none"> ◆ The volume is layered, and concatenates several underlying mirror volumes. This provides improved data redundancy and reduces the impact of disk failures. This type of volume is also known as a concatenated-mirrored volume. ◆ Recovery time is quicker than with regular concatenated volumes that are mirrored. ◆ The volume can consist of disk regions that are not adjacent. ◆ The volume can span multiple disks. A spanned volume can have a greater capacity than a single disk. 	<ul style="list-style-type: none"> ◆ A concatenated mirrored volume requires at least two disks. ◆ Concatenation does not improve I/O performance. ◆ Layered volumes consist of more VxVM objects than non-layered volumes and are likely to fill up the disk group configuration database sooner. When the configuration database is full, you cannot create more volumes in the disk group.
Striped Mirrored	<ul style="list-style-type: none"> ◆ The volume is layered, and stripes several underlying mirror volumes. This provides improved data redundancy and reduces the impact of disk failures. This type of volume is also known as a striped mirrored volume. ◆ Recovery time is quicker than with regular striped volumes that are mirrored. ◆ Striping provides improved read and write performance. ◆ Striping can help to balance the I/O load from multi-user applications across multiple disks. 	<ul style="list-style-type: none"> ◆ A striped mirrored volume requires at least four disks. ◆ Layered volumes consist of more VxVM objects than non-layered volumes and are likely to fill up the disk group configuration database sooner. When the configuration database is full, you cannot create more volumes in the disk group.

Creating Volumes

Note The total volume size you indicate for the volume in the Create Volume Wizard is the *usable* size for storing the data. Depending on the layout type, the volume can take more space on the disk. The size shown in the Wizard for simple, striped, and spanned volumes is the actual volume size on the disk. A RAID-5 volume requires additional space for parity information, and a mirrored volume is a multiple of the space taken by the original plex of the volume. The Wizard provides a **Query Max Size** button to determine the size of the volume. Again, this is the usable size of the volume rather than the actual size taken by the volume on the disk.

Creating Volumes

You can create new volumes on disks under Volume Manager control. When you invoke the new volume wizard, it asks you for information on the disks to use, the required layout, size, mirrors, logging, and file system options.

▼ To create a volume

1. Click on the volume node.
2. Choose **Actions > New Volume**. The new volume wizard appears.
3. You can suppress the Welcome to New Volume Wizard page by selecting the **Do not show this page next time** checkbox.
4. Click **Next>** to continue.

Note To reinstate the Welcome page, you need to launch the wizard, click on the **<Back** button, then uncheck the checkbox.

You can create a VxVM volume (see *Creating a VxVM Volume*), or VxVM ISP volumes (see *Creating a VxVM ISP Volume*).

Creating a VxVM Volume

Note You can create a VxVM volume by first selecting the disk group; then selecting **New Volume**, or by invoking **New Volume** from the **Volume** node. If you select from the Volume node; you will then need to select a disk group. This section describes how to create a VxVM volume by first selecting the disk group. Both methods are identical once the disk group has been selected.

When you create a VxVM volume, you can either select your own disks, or let VxVM select the disks for you:

Selecting Your Own Disks

VxVM Selecting the Disks

Selecting Your Own Disks

(Refer to VxVM Selecting the Disks, if you want to VxVM to select the disks.)

Use the following procedure to create your volume in a selected disk group:

▼ To create a volume

1. Click on the disk group to be used to create the volume under **Disk Groups** under the selected host in the left pane. (The selected disk group must not have been organized using ISP.)
2. Choose **Actions > New Volume**.
3. You are now asked if you want Volume Manager to select the disks to use; or, if you want to select them manually.
4. Click the **Manually select disks to use with this Volume** radio button, click **Next>**.
5. The **Select Disks to use for Volume** screen appears. You can now:
 - ◆ Select disks to be used by VxVM when creating the volume. (Move to **Included:** area.)
 - ◆ Select disks not to be used by VxVM when creating the volume. (Move to **Excluded:** area.)
 - ◆ Specify that the volume is to be mirrored across (**Mirror Across:**) or striped across (**Stripe Across:**) controllers, trays, targets, or enclosures.

Creating Volumes

Note For ISP Volumes the layout (**Mirror Across:** and **Strip Across:**) cannot be specified.

- ◆ Specify ordered allocation (**Ordered**). Ordered allocation uses the specified storage to first concatenate disks, then to form columns, and finally to form mirrors.
6. Click **Next>**. The **Select the attributes for this volume** screen appears.
 7. Type in the Volume Name. This is a Volume Manager-specific name that is used in some VxVM commands. It is different from the volume label for the file system.
 8. Specify the volume size, or select **Maxsize**.

If you click the **Maxsize** button, a number appears in the volume size box that represents the maximum size for a spanned volume that spans all the disks in the group. You can also click on another volume layout and then click the **Maxsize** button to get the maximum size for that layout that involves unallocated space on all disks in the group.

If you choose Striped or RAID-5, then Number of Columns and Stripe Unit Size need to have an entry. Defaults are provided.

9. Select the required layout:
 - ◆ Concatenated (Concatenated and Concatenated Mirrored)
 - ◆ Striped (Striped and Striped Mirrored)
 - ◆ RAID-5
 - ◆ Mirrored
10. Select the required options

Option	Notes
Mirror Info:	<ul style="list-style-type: none"> ◆ To mirror the volume, select Mirrored. In the Total Number of Mirrors field, enter the total number of mirrors for the volume. <p>Note Concatenated mirrored volumes are mirrored by default.</p>
Initialize zero:	<ul style="list-style-type: none"> ◆ To clear the volume before enabling it for general use, select Initialize Zero.

No layered volumes:	<ul style="list-style-type: none"> To prevent the creation of a layered volume, select No Layered Volumes. In cases where a layered volume layout is appropriate, VxVM can create a layered volume when a non-layered layout is specified. This option ensures that the volume has a non-layered layout. If a layered layout is selected, this option is ignored.
Enable FastResync	<ul style="list-style-type: none"> If the Enable FastResync box is checked; then Enable Logging is an option only if the volume is mirrored.
Enable logging	<p>(The reason it is only enabled when Enable FastResync is checked is to ensure that a new style DRL is created.)</p>

- Click **Next>** to continue. You are now asked if you want to create a File System.
- If you want to create file system, select **Create a File System** (see “Creating a New File System on a Volume Set,” , Chapter 7).
If you do not want to create a file system, select **No File System**, then click **Next>**.
- A summary of your selections appears. Click **<Back** to make changes, otherwise, click **Finish**.

VxVM Selecting the Disks

(Refer to Selecting Your Own Disks, if you want to select your own disks.)

Use the following procedure to create the volume:

▼ To create a volume

- Click on the disk group to be used to create the volume under **Disk Groups** under the selected host in the left pane.
- Choose **Actions > New Volume**. You are now asked if you want Volume Manager to select the disks to use; or, if you want to select them manually.
- Select **Let Volume Manager decide which disks to use for this volume**, then click **Next>** to bring up the **Select the attributes for this volume** screen.
- Type an alphanumeric string in the **Volume Name** field. The total number of characters must be less than 32. This is a Volume Manager-specific name that is used in some VxVM commands. It is different from the volume label for the file system.
- Specify the volume size, or select **Maxsize**.

Creating Volumes

If you click on a volume type and also click the **Maxsize** button, a size appears in the volume size box that represents the maximum volume for that layout for all disks in the group.

If you choose Striped or RAID-5, Number of Columns and Stripe Unit Size need to have an entry. Defaults are provided.

6. Select the required layout:
 - ◆ Concatenated (Concatenated and Concatenated Mirrored)
 - ◆ Striped (Striped and Striped Mirrored)
 - ◆ RAID-5
 - ◆ Mirrored
7. Select the required options

Option	Notes
Mirror Info:	To mirror the volume, select Mirrored . In the Total Number of Mirrors field, enter the total number of mirrors for the volume. Note Concatenated mirrored volumes are mirrored by default.
Initialize zero:	To clear the volume before enabling it for general use, select Initialize Zero .
No layered volumes:	To prevent the creation of a layered volume, select No Layered Volumes . In cases where a layered volume layout is appropriate, VxVM can create a layered volume when a non-layered layout is specified. This option ensures that the volume has a non-layered layout. If a layered layout is selected, this option is ignored.
Enable FastResync Enable logging	◆ If the Enable FastResync box is checked; then Enable Logging is an option only if the volume is mirrored. (The reason it is only enabled when Enable FastResync is checked is to ensure that a new style DRL is created.)

8. Click **Next>** to continue. You are now asked if you want to create a File System.
9. If you want to create file system, select **Create a File System** and proceed to **Creating File Systems**.
If you do not want to create a file system, select **No File System**, then click **Next>**.

10. A summary of your selections appears. Click **<Back** to make changes, otherwise, click **Finish**.

Creating a VxVM ISP Volume

Use the following procedure to create your VxVM ISP volume in a selected disk group:

▼ To create a VxVM ISP volume

1. Click on the disk group to be used to create the volume under **Disk Groups** under the selected host in the left pane. The selected disk group must have been organized using ISP.
2. Choose **Actions > New Volume**. You are now asked if you want to create the volume from an existing template (Create a VxVM ISP Volume from an Existing User Template), or if you want to define a new template (Create a VxVM ISP Volume Using a New Template).

Create a VxVM ISP Volume Using a New Template

Use the following procedure to create your VxVM ISP volume using a newly created template:

▼ To create a VxVM ISP volume using a new template

1. From the Organized Disk Group (Organizing Disk Groups) that you want to create your ISP-style volume with, select **Actions > New Volume**.
2. Press the **Blank User Template** button in the Select User Template dialog, then click **Next>**. The Select Capabilities dialog appears.
3. Select the capabilities you want. (Refer to Select Capabilities, Using Intelligent Storage Provisioning.) Click **Next>**. The Rule configuration dialog appears.
4. Select the rules you want. (Refer to Set Configuration Rules, Using Intelligent Storage Provisioning.) Click **Next>**. The Select Disks dialog appears.
5. You can either select the disks yourself, or let VxVM select the disks for you.
6. If you want to select the disks yourself, select the disks which can be used for the new volume (move to **Include**). You can also select those disks you want to exclude from creating the new volume (move to **Exclude**).

Creating Volumes

7. Click **Next>**. The Volume Attributes wizard appears.
8. Type an alphanumeric string in the **Volume Name** field. The total number of characters must be less than 32. This is a Volume Manager-specific name that is used in some Volume Manager commands. It is different from the volume label for the file system.
9. Specify the volume size, or select **Maxsize**.

Note If you select **Maxsize**, a size appears in the volume size box that represents the maximum value for the selected capabilities and rules.

10. Click **Next>**. You are now asked if you want to create a File System.
11. If you want to create file system, select **Create a File System** and proceed to **Creating File Systems**.
If you do not want to create a file system, select **No File System**, then click **Next>**.
12. Click on **Finish** when you are satisfied with all the settings displayed on the summary page. Use the **<Back** button to make any changes to the earlier pages.

Create a VxVM ISP Volume from an Existing User Template

If you choose to use the existing configuration (user template), then you have the option to skip the capability and rules page, and you are taken directly to the disk selection page by selecting the **Skip Wizard Pages**. (The volume attributes screen is always shown.)

Use the following procedure to create your VxVM ISP volume:

▼ To create a VxVM ISP volume using an existing template

1. From the Organized Disk Group (Organizing Disk Groups) that you want to create your ISP-style volume with, select **Actions > New Volume**.
2. Select the **Existing User Templates** button in the Select User Template dialog, then click **Next>**.
3. Select a template listed under **Available User Templates** column.
4. If you do not want to modify the capabilities and rules for the selected existing template, select the **Skip Capabilities** check box, then click **Next>**. The **Select Disks** dialog box appears.

5. To manually select the disks, move the disks to the **Include** or **Exclude** areas. Click **Next>**. The **Volume Attributes** screen appears.
6. Type an alphanumeric string in the **Volume Name** field. The total number of characters must be less than 32. This is a Volume Manager-specific name that is used in some Volume Manager commands. It is different from the volume label for the file system.
7. Specify the volume size, or select **Maxsize**.

Note If you select **Maxsize**, a size appears in the volume size box that represents the maximum value for the selected capabilities and rules.

8. Click **Next>** . You are now asked if you want to create a File System.
9. If you want to create file system, select **Create a File System** and proceed to **Creating File Systems**.
If you do not want to create a file system, select **No File System**, then click **Next>**.
10. A summary of your selections appears. Click **<Back** to make changes, otherwise, click **Finish**.

Concatenated

A concatenated volume consists of one or more regions of the specified disks. You have the option of placing a file system on the new volume or mirroring the volume. You can create a regular concatenated volume or a concatenated mirrored volume. A concatenated mirrored volume is layered and mirrored.

Layout:	Choose Concatenated or Concatenated Mirrored for the volume layout.
Options:	<ul style="list-style-type: none"> ◆ To mirror the volume, select Mirrored. In the Total Number of Mirrors field, enter the total number of mirrors for the volume. <p>Note Concatenated mirrored volumes are mirrored by default.</p> <ul style="list-style-type: none"> ◆ To enable logging for a mirrored volume, select Enable Logging. ◆ To clear the volume before enabling it for general use, select Initialize Zero. ◆ To prevent the creation of a layered volume, select No Layered Volumes. In cases where a layered volume layout is appropriate, VxVM can create a layered volume when a non-layered layout is specified. This option ensures that the volume has a non-layered layout. If a layered layout is selected, this option is ignored.

Notes:
<ul style="list-style-type: none"> ◆ If no disks are assigned, VxVM uses available space on disks in the selected disk group. ◆ The data in a concatenated volume is not protected against disk failure unless the volume is mirrored. At least one additional disk is required to mirror a concatenated volume. ◆ Concatenated mirrored volumes are mirrored by default, so a concatenated mirrored volume requires more disks than an unmirrored concatenated volume. ◆ If your system supports root volumes, then it should be noted that you cannot use a concatenated mirrored volume for a <code>root</code> volume. ◆ You cannot use a concatenated mirrored volume for a <code>swap</code> volume. ◆ If a file system is placed on the volume and the Mount at Boot option is selected, the file system table file automatically updates. ◆ The Enable FastResync option is available only for VxVM 4.0 data groups, and if the FastResync license is present.

Striped

A striped volume spreads across two or more disks. You have the option of placing a file system on the new volume or mirroring the volume. You can create a regular striped volume or a striped mirrored volume. A striped mirrored volume is layered and mirrored.

Layout:	<p>Choose Striped or Striped Mirrored for the volume layout.</p> <p>In the Number of Columns field, specify the number of columns (disks) across which the volume should be striped. A striped volume requires at least two disks.</p> <p>Specify a stripe unit size (sectors, kilobytes, megabytes, or gigabytes) in the drop-down menu.</p>
Options:	<ul style="list-style-type: none"> ◆ To mirror the volume, select Mirrored. In the Total Number of Mirrors field, enter the total number of mirrors for the volume. <p>Note Striped mirrored volumes are mirrored by default.</p> <p>To enable logging for a mirrored volume, select Enable Logging.</p> <ul style="list-style-type: none"> ◆ To clear the volume before enabling it for general use, select Initialize Zero. ◆ To prevent the creation of a layered volume, select No Layered Volumes. In cases where a layered volume layout is appropriate, VxVM can create a layered volume when a non-layered layout is specified. This option ensures that the volume has a non-layered layout. If a layered layout is selected, this option is ignored.

Notes:

- ◆ If no disks are assigned, VxVM uses available space on disks in the selected disk group.
- ◆ The data in a striped volume is not protected against disk failure unless the volume is mirrored. At least one additional disk is required to mirror a striped volume.
- ◆ Striped mirrored volumes are mirrored by default, so a striped mirrored volume requires more disks than an unmirrored striped volume.
- ◆ If your system supports root volumes, then it should be noted that you cannot use a striped mirrored volume for a `root` volume.
- ◆ You cannot use a striped mirrored volume for a `swap` volume.
- ◆ If a file system is placed on the volume and the **Mount at Boot** option is selected, the file system table file is automatically updated.
- ◆ The **Enable FastResync** option is available only for VxVM 4.0 data groups, and if the FMR3 license is present.

Creating Volumes

RAID-5

A RAID-5 volume spreads across two or more disks, and contains parity information. You have the option of placing a file system on the new volume.

Layout:	<p>Choose RAID-5 for the volume layout.</p> <p>In the Number of Columns field, specify the number of columns (disks) across which the volume should be striped. A RAID-5 volume requires at least three disks (or four disks if logging is enabled).</p> <p>To specify a stripe unit size other than the default, enter the stripe unit size in the Stripe Unit Size field. To specify a size unit, attach <i>s</i> (sectors), <i>k</i> (kilobytes), <i>m</i> (megabytes), or <i>g</i> (gigabytes) to the size.</p>
Options:	<ul style="list-style-type: none"> ◆ To enable RAID-5 logging for the volume, select Enable Logging. VxVM adds an appropriate number of logs to the volume. ◆ To clear the volume before enabling it for general use, select Initialize Zero.

Notes:

- ◆ A RAID-5 volume requires at least three disks. If logging is enabled, a RAID-5 volume requires at least four disks.
- ◆ RAID-5 volumes cannot be mirrored.
- ◆ For a RAID-5 volume, the size specified in the Size field is the usable space in the volume. Volume Manager allocates additional space for the volume parity information. The disks across which the RAID-5 volume is striped must, therefore, contain additional free space for the volume parity information.
- ◆ If no disks are assigned, VxVM uses available space on disks in the selected disk group.
- ◆ If a file system is placed on the volume and the **Mount at Boot** option is selected, the file system table file is automatically updated.
- ◆ The **Enable FastResync** option is available only for VxVM 4.0 data groups, and if the FMR3 license is present.

Mirrored

A concatenated or striped volume is mirrored (copied) on at least one other disk. If one of the disks fails, data on the volume can be accessed from the surviving disks. Concatenated mirrored and striped mirrored volumes are layered and mirrored by default.

Layout:	<p>Choose Concatenated, Striped, Concatenated Mirrored, or Striped Mirrored for the volume layout.</p> <p>For a striped or striped mirrored volume, provide the following information:</p> <ul style="list-style-type: none"> • In the Number of Columns field, specify the number of columns (disks) across which the volume should be striped. A striped volume requires at least two disks. • To specify a stripe unit size other than the default, enter the stripe unit size in the Stripe Unit Size field. To specify a size unit, attach an s (sectors), k (kilobytes), m (megabytes), or g (gigabytes) to the size.
Mirror Info:	<p>Select Mirrored. In the Total Number of Mirrors field, enter the total number of mirrors for the volume.</p>
Options:	<ul style="list-style-type: none"> • To enable Dirty Region Logging (DRL) for the volume, select Enable Logging. VxVM adds an appropriate number of logs to the volume. • To clear the volume before enabling it for general use, select Initialize Zero. • To prevent the creation of a layered volume, select No Layered Volumes. In cases where a layered volume layout is appropriate, VxVM can create a layered volume when a non-layered layout is specified. This option ensures that the volume has a non-layered layout. If a layered layout is selected, this option is ignored.

Notes:

- A mirrored volume requires at least two disks. If logging is enabled, the volume requires at least one additional disk.
- Only concatenated or striped volumes can be mirrored.
- The Prepare for Instant Snapshot option is available only for VxVM 4.0 data groups, and if the

Resizing a Volume

You can change the size of a VxVM volume. If users require more space on a volume, you can use this procedure to increase the size of the volume. If a volume contains unused space that you need to use elsewhere, you can use this procedure to shrink the volume. If the volume contains a file system, this procedure also resizes the file system.

▼ To resize a VxVM volume

1. Select the volume to be resized.
2. Choose **Actions > Resize Volume**. The Resize Volume dialog appears, displaying the name of the volume and its current size.
3. You can find out the maximum size of the volume by selecting **Max Size**.
4. Choose one of the following actions:
 - ◆ Specify the size which the volume should grow (in the **Add By:** field). The new size is displayed in **New volume size:** field.
 - ◆ Specify the size which the volume should reduce (in the **Subtract By:** field). The new size is displayed in **New volume size:** field.
 - ◆ Specify the required size in the **New volume size:** field.
 - ◆ Specify that the volume should grow to the maximum size available, by selecting **Max Size**.
5. You can force the resize, if the size is being reduced and the volume is active, by selecting the **Force** checkbox.
6. To use a specific disk for the additional space, click **Manually select disks for use by this volume**, then use the buttons for adding and removing disks to move the disk or disks you want assigned as a destination disk to the right pane of the window.

7. If you select the disks manually, you can specify the layout, select **Mirror Across** or **Stripe Across**; then select the device (Controller, Tray, Target, Enclosure). Click **OK**.

Notes:

- ◆ The **Resize Volume** task requires either the desired size or the amount of space to add to or subtract from the volume size.
- ◆ A volume containing a file system other than VxFS cannot be reduced in size.
- ◆ Shrinking a volume with an unmounted file system can result in data loss, so a volume containing an unmounted file system cannot be reduced in size.
- ◆ When the size of volume is reduced, the resulting extra space is freed for use.
- ◆ When the volume size is increased, sufficient disk space must be available. When increasing the size of a volume, VxVM assigns the necessary new space from available disks.
- ◆ You cannot **Mirror Across** or **Stripe Across** ISP Volumes when you choose to select the disks manually.

Resizing and Deleting Multiple VxVM ISP Volumes

To resize or delete multiple volumes, do the following:

1. Select the volumes to be resized or deleted
2. Right click on the selected items
3. Select and perform the appropriate action

The changes will be applied to the selected volumes.

- ❖ In case of error

If there is an error on any of the selected volumes, the changes are not applied to this volume and an error message specific to the volume is displayed. The changes are successful on the rest of the selected volumes.

Renaming a Volume

The **Rename Volume** procedure changes the name of a volume. If the volume has a file system, this procedure also makes the necessary changes to the file system table file and allows you to specify a new mount point for the file system.

Renaming a Volume

▼ To rename a volume

1. Select the volume to be renamed.
2. Choose **Actions > Rename Volume**.
3. Complete the Rename Volume dialog box as follows:

Volume Name:	If the correct volume name does not already display in this field, enter the volume name.
New Name:	Enter the new name for the volume.

After you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The **Rename Volume** task requires the current volume name and a new volume name.
- ◆ The new volume name must be unique within the disk group.
- ◆ If the volume contains a file system, the file system table file is automatically updated (if necessary).
- ◆ The Rename Volume task requires **umount** file system if a file system is mounted.

Transforming an ISP Volume

Use Transform Volume option to modify a volume's capabilities and rules. Select and right click the volume that you want to modify. Select the Transform Volume option.

The wizard that is displayed guides you through the **Volume Attributes, Capabilities, Rules, and Disk Selection** pages. Change the volumes intent to suit your current requirements.

▼ To Transform an ISP Volume

1. Select the ISP volume to transform.
2. Choose **Actions > Transform Volume**. The Transform Volume Wizard appears. Click **Next**.
3. The **Volume Attribute** dialog appears. This dialog displays the volume to be transformed and the size of the volume. In this dialog a comment about the volume can be added.
4. Click **Next**. You are now asked if you want select existing template (Create a VxVM ISP Volume from an Existing User Template), or if you want to define a new template (Create a VxVM ISP Volume Using a New Template). Once you have chosen the template you wish to use, click **Next**.
5. Select the capabilities you want. (Refer to Select Capabilities, Using Intelligent Storage Provisioning.) Click **Next**.
6. If asked, select the rules you want. (Refer to Set Configuration Rules, Using Intelligent Storage Provisioning.) Click **Next**.
7. A summary of your selections appears. Click **<Back** to make changes, otherwise, click **Finish**.

Adding and Removing Columns in a VxVM ISP Volume

You can change the number of columns in a VxVM ISP volume using the **Add/Remove Column** menu item.

▼ To change the number of columns in a VxVM ISP volume

1. Select the volume to be changed.
2. Choose **Actions > Add/Remove Column**. The Add/Remove Column dialog appears. This displays the number of columns and the stripe size on the selected volume. Complete options for changing the layout as follows:

Columns	Increase or decrease the number of columns to Add or Remove columns for the volume.
Stripe Width	Change the stripe width of the columns in the Stripe Width field.

Changing the VxVM ISP Column Count

The volume **Change Column Count** procedure changes the column count and strip unit size of a volume without interfering with the performance of a volume.

▼ To change an ISP column count

1. Select the volume to be changed to a different column count.
2. Choose **Actions > Change Column Count**. The Change Column Count dialog appears. This displays the number of columns and the stripe size on the selected volume. Complete options for changing the layout as follows:

Columns	Increase or decrease the number of columns to Add or Remove columns for the volume.
Stripe Unit Size	Change the stripe unit size of the columns in the Stripe Unit Size field.

3. Click **OK** to apply the changes.

Setting Volume Usage

The volume **Set Volume Usage** procedure enables you to set the read policy for a volume.

▼ To set volume usage

1. Select the volume.
2. Choose **Actions > Set Volume Usage**.

Changing a VxVM Volume Layout

The volume **Change Layout** procedure changes the layout of a VxVM volume without interrupting access to data. You can use this task to change the redundancy or performance characteristics of an existing volume.

You can view the progress of the relayout task through the **Relayout Status Monitor** window. You can also use the **Relayout Status Monitor** window to pause, abort, continue, or reverse the relayout task. The **Relayout Status Monitor** appears automatically when a relayout task begins.

▼ To change a volume layout

1. Select the volume to be changed to a different layout.
2. Choose **Actions > Change Layout**.
3. Complete the **Change Volume Layout** dialog box as follows:

Volume Name:	If the correct volume name does not already display in this field, enter the volume name.
Layout:	<p>Choose the new volume layout (Concatenated, Striped, RAID-5, Concatenated Mirrored, or Striped Mirrored.)</p> <p>For a striped, RAID-5, or striped mirrored volume, provide the following information:</p> <ul style="list-style-type: none"> ◆ In the Number of Columns field, specify the number of columns (disks) across which the volume should be striped. ◆ The stripe unit size (sectors) in the Stripe Unit Size field.

Changing a VxVM Volume Layout

Options:	<p>Click Show Options to access the optional settings for this task.</p> <ul style="list-style-type: none"> ◆ To retain the original volume size when the volume layout changes, click Retain volume size at completion. ◆ To specify the size of the pieces of data that are copied to temporary space during the volume relayout, enter the size in the Temp space size field. ◆ To specify additional disk space to be used for the new volume layout (if needed), specify the disk in the disk(s) field. ◆ To specify the temporary disk space to be used during the volume layout change, specify the disk in the temp disk(s) field. ◆ If the volume contains plexes with different layouts, specify the plex to be changed to the new layout in the Target Plex field.
----------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

4. After you have provided all necessary information in the dialog box, click **OK**. The Relayout Status Monitor window appears. This window provides you with the following information:

Volume Name:	The name of the volume that is undergoing a layout change.
Initial Layout:	The original layout of the volume.
Desired Layout:	The new layout for the volume.
Status:	The status of the relayout task.
% Complete:	The progress of the relayout task.

If you need to pause, abort, or continue a layout task after you close the **Relayout Status Monitor** window, you can access the relayout task through the popup menu at the bottom of the main window, see “**Accessing Tasks**” for more information. You can use the popup menu to abort, pause, or resume the task.

Notes:

- ◆ The **Volume Change Layout** task requires a volume name and a different volume layout.
- ◆ Volumes remain accessible during online relayout.
- ◆ This task may take a long time, depending on the volume size and other factors. The **Relayout Status Monitor** shows the progress of the layout change and allows you to abort, pause, or reverse the task.
- ◆ If no disk space is specified, VxVM uses available space on disks in the current disk group.
- ◆ If the temporary space size is not specified, VxVM uses an appropriate size. Specifying a larger temporary space size speeds up the layout change process because larger pieces of data are copied at a time. If the specified temporary space size is too small, VxVM uses a larger size.
- ◆ If all of the plexes in the volume have identical layouts, VxVM changes all plexes (except log plexes) to the new layout. If the volume contains plexes with different layouts, a target plex must be specified. When a target plex is specified, VxVM changes the layout of the target plex and leaves the other plexes in the volume alone.
- ◆ Log plex layouts do not change. VxVM removes and/or replaces log plexes, depending on the new layout.
- ◆ If the volume length changes as a result of the layout change and the volume contains a file system, the file system length is automatically adjusted (if possible).
- ◆ A volume cannot undergo multiple relayouts at the same time.
- ◆ This task cannot be used to change the layout of a volume with a sparse plex.
- ◆ This task may fail for volumes that were not created by VEA or the `vxassist` command.
- ◆ If the relayout is reversed, the volume returns to its original layout. However, data in the volume may not return to its original storage location.
- ◆ During a relayout, many other tasks are unavailable for the volume.
- ◆ A mirrored or striped volume requires at least two disks.
- ◆ A RAID-5 volume requires at least three disks.

Adding a Mirror to a Volume

The volume **Add Mirror** procedure creates a mirror (copy) of a volume on a disk that is not already being used by the volume. If one of the disks fails, the volume data can be accessed from the surviving disks. A volume can have multiple mirrors.

(To create a new mirrored volume, use the **New Volume** procedure. You select a mirrored layout when you create the volume. See "VxVM Volumes.")

Note The mirroring process is disabled if the volume cannot be mirrored or if there is no unallocated space on another disk large enough to mirror the volume. A RAID-5 volume cannot be mirrored.

▼ To add a mirror to an existing volume

1. Select the volume to be mirrored. You can view the mirrors on a volume by selecting the **Mirrors** tab from the volume's detail view.
2. Choose **Actions > Mirror > Add**.
3. Complete the Add Mirror dialog box as follows:

Options:	<ul style="list-style-type: none"> ◆ To add more than one mirror, enter the number of mirrors in the Number of Mirrors to Add field. ◆ To place the mirror on a specific disk, select Manually select disks for use by this volume, and assign the destination disks for the mirrors. ◆ To specify the layout for the mirror, select Mirror Across or Stripe Across. For a striped mirror, specify the number of columns and stripe unit size. <p>Note The layout does not need to be specified for an ISP Volume. Also, Option 3 (Mirror Across/Stripe Across) is not applicable to ISP volumes.</p> <ul style="list-style-type: none"> ◆ To let VxVM assign the disks, select Let Volume Manager decide what disks to use for this volume.
----------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

After you have provided all necessary information in the dialog box, click **OK**. Adding a mirror requires resynchronization, so this task may take some time.

Notes:

- ◆ Once mirrored, the data in the volume is redundant. If a disk fails, the data remains available on the surviving mirror (on another disk).
- ◆ The new mirror cannot be created on a disk that already contains a copy of the volume. Each mirror must reside on a separate disk.
- ◆ Sufficient disk space must be available to accommodate the additional mirrors.
- ◆ Only disks in the same disk group as the volume can be used to create the new mirror.
- ◆ If no disks are assigned, VxVM uses available disk space to create the mirror.
- ◆ A RAID-5 volume cannot be mirrored.
- ◆ A volume can contain up to 32 plexes (mirrors).
- ◆ If you create more than one mirror at a time, you may see inconsistent information on the progress bar. Also, the generation of multiple mirrors does affect system resources. After creating a mirror, you may want to wait until it has finished generating before creating another mirror.

Adding a Log to a Volume

The volume **Add Log** procedure adds the appropriate type of log to a mirrored or RAID-5 volume.

When used to add a Dirty Region Logging (DRL) log to a mirrored volume, DRL uses this log to track the regions of the volume that change due to I/O writes. If a system failure occurs, DRL uses the information in the log to recover only the portions of the volume that need recovery. This speeds up recovery time for mirrored volumes.

For DRL to be in effect, a mirrored volume must have at least one DRL log. You can create additional DRL logs (on different disks) to mirror the DRL information.

When used to add a log to a RAID-5 volume, the log speeds up the resynchronization time for RAID-5 volumes after a system failure. A RAID-5 log maintains a copy of the data and parity being written to the volume at any given time. If a system failure occurs, VxVM can “replay” the RAID-5 log to resynchronize the volume. This copies the data and parity that was being written at the time of failure from the log to the appropriate areas of the RAID-5 volume. You can create multiple RAID-5 logs (on different disks) to mirror the log information. Ideally, each RAID-5 volume has at least two logs.

1. Select the mirrored volume or RAID-5 volume to contain the log.
2. Select **Actions > Log > Add**.
3. In the **Number of Logs to Add** field, select the number of logs you wish to add to this volume.

Note This field is only available if you are adding logs to an ISP Volume.

4. By default the **Auto select disks** option will be selected. To manually assign the destination disks, select the **Manually select disks** option.

Available disks will appear in the left window. To add destination disks, select them in the left window and click the **Add** button. To add all of the available disks, click the **Add All** button.

To remove destination disks, select them in the right window and then click the **Remove** button. To remove all of the destination disks, click the **Remove All** button.

5. Click **Apply** to create the Log then **OK**.

Notes:

- ◆ When a log is added to a mirrored volume, DRL is activated for that volume.
- ◆ You cannot add a DRL log to an ISP Volume. ISP Volumes can only contain DCO logs.
- ◆ RAID5 volume users can add either a DCO log or a RAID5 log.

Stopping and Recovering a Volume

The **Stop Volume** procedure disables a volume. When a volume is stopped, it is not available for use until it is recovered. (refer to Recovering a Volume). Stopping a volume is a safe way of temporarily preventing access to a volume.

▼ To stop a volume

1. Select the volume to be stopped.
2. Choose **Actions > Stop Volume**.
3. In reply to the message “Are you sure you want to stop the volume ...,” click **OK**. The volume state changes to STOPPED.

Note:

- ◆ A volume cannot be stopped if it is in use (open) or it has a mounted file system.

Recovering a Volume

The **Recover Volume** procedure performs any necessary volume recovery operations on a volume. The recovery operations depend on the type of volume and include starting disabled volumes, resynchronizing mirrors in mirrored volumes, and resynchronizing parity in RAID-5 volumes. After successful recovery, the volume is again available for use.

▼ To recover a failed volume

1. Select the volume to be recovered.
2. Choose **Actions > Recover Volume**.
3. In reply to the message “Are you sure you want to recover volume ...,” click **OK**.

Note:

In some cases, recovery is not possible. If the volume recovery fails, you can attempt to restore the volume from backup.

Preparing to Restore a Volume From Backup

The **Volume Prepare Restore** procedure stops the volume, sets the volume to an uninitialized state, and restarts the volume (without resynchronizing the volume mirrors). The volume can then be reloaded from backup (refer also to *Merging a Snapshot Volume (Volume Snapback)*).

This procedure is intended for raw image backups and is not appropriate for file system backups. This procedure is useful for disaster recovery. If the data on a volume is corrupted and you know that you need to restore the volume from backup, you can use this procedure to prepare the volume for restoration.

▼ To prepare a volume for a restore from backup

1. Select the volume to be restored from backup.
2. Choose **Volumes > Restore Volume**. When asked, "Are you sure you want to restore this volume?," click **OK**.

Note:

If the volume contains a mounted file system, the file system must be unmounted before this task can proceed. This task does not remount the file system.

Removing a Mirror From a Volume

Caution Removing a mirror can result in loss of data redundancy.

The **Mirror Remove** procedure removes a mirror from its volume. This procedure breaks the link between the mirror and its volume and returns the mirror disk space to the free space pool for reuse. You can remove an extra mirror from a volume to make the underlying disk space available for use elsewhere. However, this may leave the volume unmirrored and unprotected against disk failure.

▼ **To remove a mirror from a volume**

1. Select the volume that contains the mirror to be removed.
2. Choose **Actions > Mirror > Remove**.
3. Complete the Remove Mirror dialog box as follows:

Volume Name:	If the correct volume name does not already display in this field, enter the volume name or select the volume from the drop-down menu.
Removal Method:	Specify the mirrors and/or number of mirrors to be removed: <ul style="list-style-type: none"> ◆ To specify how many mirrors should be removed, select By Quantity/Disk. Enter the number of mirrors to be removed in the Mirror quantity field. ◆ To specify the disks on which mirrors should be preserved, select By quantity/disk. ◆ Specify the names of disks or mirrors to be removed from the list of available disks or mirrors.

After you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The **Mirror Remove** task requires a volume name. This task also requires the number of mirrors to remove, a disk name, or a mirror name.
- ◆ If a volume only has two mirrors and one mirror is removed, the volume is no longer redundant.
- ◆ The last mirror (plex) cannot be removed from a volume.
- ◆ When a mirror is removed, the space occupied by that mirror is returned to the free space pool.

Removing a Log From a Volume

Removing a Log From a Volume

The **Log Remove** procedure removes a log from a mirrored, ISP-type, or RAID-5 volume. For a mirrored volume, this removes a DRL log from the volume. For an ISP-type volume, this removes the DCO log from the volume. For a RAID-5 volume, this removes a RAID-5 log from the volume.

▼ To remove a log from a volume

1. Select the volume that contains the RAID-5, DRL, or DCO log to be removed.
2. Select **Actions > Log > Remove**.
3. Complete the Remove Log dialog box as follows:

Volume Name:	If the correct volume name does not already display in this field, select the volume name from the drop-down menu.
Removal Method:	Specify the logs and/or number of logs to be removed: <ul style="list-style-type: none"> ◆ To specify how many logs should be removed, select By Quantity/Disk. Enter the number of logs to be removed in the Log quantity field. ◆ To specify the disks on which logs should be preserved, select By Quantity/Disk. Select the disk names from the list of available disks. ◆ To specify the name of the log to be removed, select Log, then select the log names from the list of available logs.

After you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ If the only log is removed from a volume, then logging (either dirty region logging, DCO logging, or RAID-5 logging) is no longer in effect for that volume. When logging is disabled, recovery time increases.

Deleting a Volume

Caution Deleting a volume can result in data loss.

This procedure permanently removes a volume and all of its data. You should only remove a volume if you are sure that you do not need the data in the volume (or the data is backed up elsewhere). You can remove a volume to make the underlying disk space available for use elsewhere.

▼ To delete a volume

1. You can select the volume to be deleted in one of two ways:

Select the disk group containing the volume to be deleted. The disk group table is displayed in the right pane. Click on the **Volumes** tag. The volumes in the disk group are displayed.

or

Select the **Volumes** node.

2. Select the volume to be deleted.
3. Choose **Actions > Delete Volume**.
4. Respond **Yes** when the Delete Volume confirmation dialog appears.

Notes:

- ◆ Removing a volume destroys all of the data in that volume.
- ◆ When a volume is removed, the space it occupied is available as free space.

Attaching Mirrors

The **Attach Mirror** procedure attaches a mirror to a volume. This involves copying data from an active mirror on the volume to the mirror.

▼ **To attach a mirror**

1. Select the volume containing the mirror. The volume table appears in the right pane.
2. Click on the **Mirrors** tab in the right pane. The list of mirrors for that volume appears.
3. Select mirror to be attached. Click **Attach Mirror**.

Detaching a Mirror

The **Detach Mirror** procedure dissociates a mirror from its volume. This disconnects the mirror from its volume and reduces volume data redundancy.

▼ **To detach a mirror**

1. Select the volume containing the mirror. The volume table appears in the right pane.
2. Click on the **Mirrors** tab in the right pane. The list of mirrors for that volume appears.
3. Select mirror to be detached. Click **Detach Mirror**.

Moving, Splitting, and Joining Subdisks

The **Subdisk Move**, **Subdisk Split**, and **Subdisk Join**, procedures enable you to move subdisks to other locations within the group, split subdisks, and join them back together. The flexibility of moving subdisks, splitting them, and joining them lets you make best use of your disk space.

- ◆ Moving a Subdisk
- ◆ Splitting a Subdisk
- ◆ Joining a Subdisk

After selecting the disk group, you can then use the Disk View window (**Actions > Disk View**) or the Volume to Disk Mapping window (**Actions > Disk/Volume Map**) to view the subdisks and gaps on disks in a disk group. You can use the Volume Layout Details window (**Actions > Layout View**) to view subdisks in a specific volume.

Moving a Subdisk

Note You are not allowed to move a subdisk that is part of a boot or system volume.

The **Subdisk Move** procedure moves the contents of a volume subdisk from one disk to another. This is useful for moving portions of a volume to a different disk for improved performance.

If disk activities are heavily concentrated on one or a small number of disks in the storage subsystem, it may create bottlenecks. You can use the **Subdisk Move** and possibly the **Subdisk Split** procedure to spread out disk accesses more evenly across all the disks to balance the load.

You can use the **Subdisk Move** procedure to move all subdisks to healthier disks. The benefit of moving subdisks instead of copying the volumes is that you need only enough space on the receiving disks for the subdisks on the one failed disk, not for entire volumes that may span multiple disks. Another advantage is that there is no interruption in I/O.

Moving subdisks provides you with a flexible means of making adjustments in your storage system while it is up and running.

In the Disk View window, you can move a subdisk by placing the pointer over the subdisk and then pressing and holding the mouse button to drag the subdisk to another location.

To display the subdisk in the right pane, click on its volume in the tree view, and then click on the Subdisks tab in the right pane. The subdisks for that volume will display. You can use the Sub-disk Properties menu selection to further identify the characteristics of a volume's subdisks.

Moving, Splitting, and Joining Subdisks

▼ To perform a subdisk move operation:

1. Select the volume containing the subdisk to be moved.
2. Choose **Actions > Disk View**. Then select the **Subdisks** tab.
3. Select the subdisk to be moved.
4. Choose **Actions > Move Subdisk**. The Move Subdisk dialog is displayed.
5. Choose **Auto assign destination disks**, if you want VxVM to assign a destination disk for you. Choose **Manually assign destination disks** if you want to select the destination disk.
6. If you select manual assignment of the destination disk, you need to select the destination disks. You also need to set the disk offset and its units (Sectors, KB, MB, GB, or TB).
7. Click **OK** to start the move operation.

Note VEA automatically checks the state of the data before beginning the subdisk move. If the data is degraded, the subdisk can still be moved. After moving, you will need to reactivate the disk to clear the degraded setting.

8. Click **Finish** when the final screen appears to complete the command.

Notes:

- ◆ Subdisks are typically moved to reorganize disk space.
- ◆ The disk space occupied by the original subdisk is returned to the free space pool.
- ◆ A subdisk in a mirrored, striped, or RAID-5 volume should not be moved to a disk that already contains a copy or part of that volume.
- ◆ If the **Subdisk Move** task fails and leaves some unused subdisks (that is, subdisks that are not associated with a volume) on the system, you can use the **Subdisk Remove** task to free the space occupied by the unused subdisks.

Sub-disk Properties

Property	Description
Name	Subdisks on a disk are assigned numbers in a sequence, such as 1-01, 1-02, 1-03 or 12-01, 12-02, 12-03. The subdisk number is an internal number generated by the program. There is no particular significance to the one or two digits before the dash except that, for any one disk, the first one or two numbers identifying its subdisks must be the same. The numbers after the dash represent a number for the subdisk within the disk. Thus, if the subdisk was the first one created for a particular disk, it would have 01 after the dash, the second disk created would have 02 after the dash, and so on.
Size	This field displays the size of the subdisk.
Column	Column displays the column that the subdisk occupies in the volume. Striped and RAID-5 volumes have multiple columns.
Disk Offset	This is the offset where the subdisk starts on the disk. The reason this is included in the Properties window is because for a subdisk to be split, it must be contiguous to the subdisk next to it. By looking at the offsets, you can determine which subdisks are contiguous.
Plex Offset	This is the offset of a subdisk within a plex. It shows the distance from the start of a plex to the start of a specific subdisk. The first subdisk within a plex will have an offset of "0." The second subdisk will be offset by the size of the first subdisk.

Splitting a Subdisk

Use the **Subdisk Split** procedure to split a subdisk. After the subdisk is split, the resulting two subdisks will reside on the same section of the same disk. You can then use the **Subdisk Move** procedure to move one or both subdisks to other disks. You may want to split and move a subdisk to make the best use of your disk space. A split subdisk can be joined back together again with the **Subdisk Join** command.

▼ To split a subdisk into multiple subdisks

1. Select the volume containing the subdisk to be split.
2. Choose **Actions > Disk View**. Then select the **Subdisks** tab.
3. Select the subdisk to be split.
4. Select **Actions > Split Subdisk**.

Moving, Splitting, and Joining Subdisks

5. In the next screen that comes up, specify the subdisk size for the first of the two subdisks and Click **Next**.

The remaining space left over will be assigned to the second subdisk. For example, if the original subdisk was 100 MB and you assigned 40 MB to the first subdisk, the other subdisk will be 60 MB.

6. The final screen of the Subdisk Split wizard appears. Click **Finish** to complete the command.

After a few moments, two subdisks will appear in the right pane of the Volume Manager GUI.

Notes:

- ◆ The name of the first subdisk remains the same as the selected subdisk. Other subdisks are automatically named by VxVM.
- ◆ The new, smaller subdisks occupy the same regions of the disk that the original subdisk occupied.
- ◆ The original subdisk must contain a sufficient number of sectors for the specified split to work.
- ◆ A Dirty Region Logging log subdisk cannot be split.

Joining a Subdisk

The **Subdisk Join** procedure joins two or more subdisks together to form a single larger subdisk. Subdisks can only be joined together if they belong to the same volume and occupy adjacent regions of the same disk and mirror. You can join two subdisks that had been one subdisk but that were split by the **Subdisk Split** command. You cannot split a subdisk and join it back to another subdisk that it was not split from. There also has to be room on the disk for the two subdisks.

▼ To join subdisks

1. Select the volume containing the subdisks to be joined.
2. Choose **Actions > Disk View**. Then select the **Subdisks** tab.
3. Select one of the subdisks to be joined.
4. Select **Actions > Join Subdisk**.

5. The welcome screen of the Subdisk Join wizard appears. Click **Next** to continue.
6. In the screen that comes up next, select the subdisks to be joined together and click **Next**.
7. The final screen of the Subdisk Join wizard appears. Click **Finish** to complete the operation.

Within a short time, the two subdisks will be replaced in the right pane of the Volume Manager GUI with one subdisk.

After you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The **Subdisk Join** task requires two or more existing subdisk names and a new (joined) subdisk name.
- ◆ Only contiguous subdisks on the same disk can be joined.
- ◆ The subdisks to be joined must belong to the same volume and mirror.
- ◆ For a striped volume, the subdisks must be in the same column.
- ◆ The joined subdisk name can be the name of one of the subdisks being joined.

Viewing Volume Properties


The Volume Properties window contains detailed information about the selected volume.

- ◆ To view the properties of a standard VxVM volume or a VxVM ISP Volume, select the volume and then choose **Properties** from the popup menu.

The following information is displayed.

Category	Description
Layout	Concatenated, striped, RAID-5, mirrored, or striped
Status	The normal status is "Healthy," which means that volume is accessible and has no known problems. Please refer to Troubleshooting for a complete description of other condition states.
Copies	Number of copies of the volume
Log	Indicates if logging is enabled. (Shows "FastResync" if enabled with instant snapshot.)
Read policy	Indicates the mode in which the device (or volume) is read.
Group name	The name of the disk group in which the volume was created.
FastResync	Indicates if FastResync is disabled, or is ready for instant snapshot
Region size	Space within the volume used for the policies enabled.
Pool name	Name of the storage pool within the disk group
Set name	Name of template set associated with the disk group
Raw device path	Location of the disk device associated with VxVM.
Block device path	Exact location of the device associated with VxVM.
Size	Size of volume. Note Any changes made to the units field will not be preserved on leaving the dialog.

Category	Description
Capabilities	A capability is a simple name used to indicate what the capability means for an ISP volume. Storage is allocated according to your request. For example, you can ask for a volume to be Performant and Reliable. These are capabilities have parameters such as the number of striped columns or number of mirrors.
Rules:	Rules specify the storage allocation criteria for ISP volumes. These define storage selection and storage layout. Rules are simple sentences written in the ISP language.
Comment	You can insert comments in this field. The text will be preserved by clicking OK . (This is the only field that is preserved on leaving the dialog.)
User	Name of the user for this version of VxVM.
Group	Name of the group that "owns" the user managing this version of VxVM.
Mode	The means by which this instance of VxVM is operating.



Viewing Volume Properties



Volume Set Tasks

5

VxVM Volume Sets tasks include:

- ◆ Creating a Volume Set
- ◆ Adding a Volume to a Volume Set
- ◆ Renaming a Volume Set
- ◆ Removing a Volume from a Volume Set
- ◆ Recovering a Volume Set
- ◆ Deleting a Volume Set

Volume Set Task Roadmap

The following table provides a summary of common VxVM volume set tasks.

To perform this task ...	Use this procedure ...
Create new volume sets	Creating a Volume Set
Change a volume name	Renaming a Volume Set
Remove a volume to a volume set	Removing a Volume from a Volume Set
Remove a volume set	Deleting a Volume Set
Renaming a volume set	Renaming a Volume Set
Add a volume to a volume set	Adding a Volume to a Volume Set

VxVM Volume Sets

The volume set feature of VxVM allows several volumes to be treated as a single object with one logical I/O interface. The following procedures are provided for maintaining volume sets:

- ◆ Displaying Volume Set Information
- ◆ Creating a Volume Set
- ◆ Adding a Volume to a Volume Set
- ◆ Removing a Volume from a Volume Set
- ◆ Renaming a Volume Set
- ◆ Deleting a Volume Set

Displaying Volume Set Information

The VxVM “Volume Sets” feature allows for several volumes to be treated as a single object with one logical I/O interface. I/O to and from “Volume Sets” is directed to the underlying volumes through the logical I/O interface. This feature only provides support for VxFS Multi-Component file systems.

A “Volume Sets” folder is displayed if you have already created volume sets on your system.

Under the “Volume” folder, then, if a volume is part of a volume set:

- ◆ The set name for a volume is shown in the volume table.
- ◆ The volume set table displays:
 - ◆ The associated disk group
 - ◆ The associated volumes

Volume sets information is also displayed with the disk group table. A Volume Sets tab appears in the disk group table when you select a disk group with at least one volume set.

The volume property sheet lists the volume sets of which a volume is a member. You can also use the volume set property sheet to record your own comments about your volume sets.

Creating a Volume Set

Use the following procedure to create a volume set:

▼ To create a volume set

1. Click on a volume to be included in the volume set.
2. Choose **Actions > Volume Set > Create Volume Set**. The Create Volume Set display appears.
3. You can now change the disk group, or continue with the selected disk group and volume.
4. Enter the name to be given to the volume set.
5. Move the volumes to be included in the volume set from the **Available volumes** column to the **Selected volumes** column.
6. Click **OK**.

Adding a Volume to a Volume Set

Use the following procedure to add a volume to an existing volume set:

▼ To add a volume

1. Click on a volume to be included in the volume set.
2. Choose **Actions > Volume Set > Add Volume to Volume Set**. The Add Volume to Volume Set display appears.
3. Select the volume set to which the volume is to be added.
4. Move the volumes to be added to the volume set from the **Available Volumes** column to the **Selected Volumes** column.
5. If a file system is present on the volume set, you can select the **Add Volume to Multi-Component File System** checkbox to also add the volume to the MDS.
6. Click **OK**.

Renaming a Volume Set

Use the following procedure to rename an existing volume set:

▼ **To rename a volume set**

1. Click on the volume set to be renamed.
2. Choose **Actions > Volume Set > Rename Volume Set**. The Rename Volume Set display appears.
3. Enter the new name for the volume set.
4. Click **OK**.

Recovering a Volume Set

The **Recover Volume Set** procedure performs any necessary volume set recovery operations on a volume set. The recovery operations depend on the type of volume set and include starting disabled volumes, resynchronizing mirrors in mirrored volumes, and resynchronizing parity in RAID-5 volumes. After successful recovery, the volume set is again available for use.

▼ **To recover a failed volume set**

1. Select the volume set to be recovered.
2. Choose **Actions > Recover Volume Set**.
3. In reply to the message "Are you sure you want to recover volume set...", click **OK**.

Note:

In some cases, recovery is not possible. If the volume set recovery fails, you can attempt to restore the volume set from backup.

Deleting a Volume Set

Use the following procedure to delete an existing volume set:

▼ To delete a volume set

1. Click on the volume set to be deleted.
2. Choose **Actions > Volume Set > Delete Volume Set**. The Delete Volume Set confirmation message appears.
3. Click **Yes** to confirm deletion.


Removing a Volume from a Volume Set

Use the following procedure to remove a volume from an existing volume set:



▼ To remove a volume

1. Click on the volume set folder, then click on the volume to be removed.
2. Choose **Actions > Volume Set > Remove Volume from Volume Set**. The Remove Volume from Volume Set display appears.
3. You can now select a new volume set, or continue with the already selected volume set.
4. Move the volumes to be removed from the volume set from the **Available Volumes** column to the **Selected Volumes** column.
5. Click **OK**.

Note A volume with index 0 cannot be removed while other volumes exist in the volume set. Removing a volume with index 0 can lead to data corruption if it is removed before all other volumes in the volume set have been removed.



Removing a Volume from a Volume Set



Using Volume Snapshots

6

VEA provides you with easy-to-use interfaces for performing tasks associated with volume snapshots:

- ◆ Using the Instant Snapshot Feature
- ◆ Creating Instant Snapshots
- ◆ Using the Full Volume Snapshot Feature
- ◆ Displaying the Snapshots for a Volume
- ◆ Displaying Cache Volume Information

Volume Snapshots Task Roadmap

The following table provides a summary of common volume snapshot tasks.

To perform this task ...	Use this procedure ...
<i>Instant Snapshot Tasks:</i>	
Enabling FastResync:	Enabling FastResync
Disabling FastResync:	Disabling FastResync
Create new instant snapshots	Using the Instant Snapshot Feature
Start Synchronization	Start Synchronization
Refreshing an instant snapshot	Refreshing an Instant Snapshot
Reattaching an instant snapshot	Reattaching an Instant Snapshot
Restoring a volume from an instant snapshots	Restoring a Volume from an Instant Snapshot
Disassociating an instant snapshot	Dissociating an Instant Snapshot
Splitting an instant snapshot hierarchy	Splitting an Instant Snapshot Hierarchy
<i>Full Snapshot Tasks:</i>	
Backing up a volume	Backing up a Volume
Preparing to restore a volume from backup	Preparing to Restore a Volume From Backup
Create a snapshot copy of a volume for backup	Using the Full Volume Snapshot Feature
Merge a snapshot copy of a volume with the original volume	Merging a Snapshot Volume (Volume Snapback)
Permanently break the association between a snapshot and its original volume	Dissociating a Snapshot Volume (Volume Snap Clear)
Abort the taking of a snapshot	Aborting a Snapshot
Attempt to recover all volumes on a disk	Recovering Volumes on a Disk (see Chapter 3)

Using the Volume Snapshot Features

VEA provides the capability for taking an image of a volume at a given point in time. Such an image is referred to as a *volume snapshot*.

Volume snapshots allow you to make backup copies of your volumes online with minimal interruption to users. You can then use the backup copies to restore data that has been lost due to disk failure, software errors or human mistakes, or to create replica volumes for purposes of report generation, application development, or testing.

Note To take a snapshot of a volume, you must have an empty disk in the volume's disk group.

Snapshot creates a snapshot image of a volume, which can be detached from the volume. The detached snapshot image can be used for backup, or other purposes, while the original volume remains in service. The snapshot can then be quickly reattached and resynchronized with the volume. The snapshot volume can also be permanently disconnected from the original volume.

You can attach a snapshot to either the original volume or the snapshot volume.

The following sections describe how to create and administer volume snapshots:

- ◆ Using the Instant Snapshot Feature describes how to create and administer third-mirror break-off and space-optimized snapshots of 4.0 (version 110) disk group volumes.
- ◆ Using the Full Volume Snapshot Feature describes how to create and administer traditional third-mirror break-off snapshots of pre-4.0 (pre-version 110) disk group volumes.

Note Please refer to the *VERITAS Volume Manager Administration Guide* for a full description of third-mirror snapshots.

Note Instant snapshots are the preferred mechanism for implementing online and off-host point-in-time copy solutions in VEA. Support for traditional third-mirror snapshots may be removed in a future release.

Taking Snapshots of Volume Sets

Please note the following limitation when creating full snapshots from volume sets:

A snapshot of a volume set must itself be a volume set. A volume set can be formed by breaking off existing plexes and must be created from a prepared volume set.

Using the Instant Snapshot Feature

The instant snapshot feature allows you to create snapshots of 4.0 (version 110) disk group volumes.

(Refer to Using the Full Volume Snapshot Feature for information on creating snapshots of pre-4.0 (pre-version 110) disk group volumes.)

Note The information in this section also applies to RAID-5 volumes converted to a special layered volume layout by the addition of a DCO (Data Change Object) and DCO volume.

The **Create** procedure creates an instant snapshot that is immediately available for making a backup. Read requests for data in the original volume at the point in time that the snapshot was taken are satisfied by reading either from a non-updated region of the original volume, or from the copy of the original contents of an updated region that have been recorded by the snapshot.

Note The information in this section also applies to creating snapshots from volume sets.

Using the FastResync Feature

If you have already enabled FastResync using a previous version of VxVM (for example, you might have used FMR2 with VxVM 3.5), you need to remove any DCOs from the disk groups to be upgraded to 4.0 (version 110).

Enabling FastResync

The **Enable FastResync** procedure gives you the option of preparing a volume for snapshot in advance of using the create instant snapshot feature. This makes a volume ready for instant snapshot and space-optimized snapshot operations. It can also prepare it for Dirty region Logging (DRL), if required.

Note If you have enabled FastResync on a previous version of VxVM, you need to remove any old DCOs from the volumes to be used for instant snapshot.

▼ To enable FastResync on a volume prior to instant snapshot

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Enable FastResync**. The Enable FastResync dialog appears.

3. If DCOs are already present as a result of enabling FastResync on a previous version of VxVM, you are now given the option of removing them proceeding. Failure to do so will result in the following message:

The removal of the previous version of FastResync has failed. Cannot proceed to using instant snapshot feature.

4. Provided that there are no existing DCOs, you can now elect the options you require:

FastResync (DCO) Mirrors:	Provide the number of DCO volumes to be created.
Region Size (KB):	Size of region (in kilobytes).
Enable Dirty Region Logging	Check this box if you require DRL.
Let Volume Manager decide which disks to use for this volume	Check this box if you do not wish to manually select disks for the volume.
Manually select disks for use by this volume	Check this box to manually select disks for use by the volume.

5. Select **OK**
6. You can now create an instant snapshot. To create instant snapshots, see “Creating Instant Snapshots” on page 136.

Creating Instant Snapshots

You can create three types of instant snapshot:

- ◆ Full sized; a full sized snapshot is a copy of the original volume. These snapshots are similar to the traditional third-mirror snapshots, but are available as soon as the plexes have been created.

Note A plex in an instant (copy-on-write) full-sized snapshot requires as much space as the original volume. If you instead make an instant space-optimized snapshot of a volume, this only requires enough storage to record the original contents of the parent volume as they are changed during the life of the snapshot.

- ◆ Break off; equivalent to the traditional third-mirror snapshot (where the snapshot plex can be detached from the original volume into a new disk group).

Note Please refer to the *VERITAS Volume Manager Administration Guide* for a full description of break-off snapshots.

- ◆ Space optimized; uses a storage cache to reduce the size of the snapshot volume.

If FastResync is not Enabled

If FastResync is not enabled, you must first go through the FastResync procedure before you can create an instant snapshot.

▼ To enable FastResync

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Create**. The Create Instant Snapshot wizard appears.
3. In the **Choose the method by which to select Disks for FastResync** dialog, you can either let VxVM select the disks for you (**Let Volume Manager decide what disks to use for FastResync**), refer to VxVM Selecting the Disk for FastResync, or select your own disks (**Manually select disks to use for FastResync**), refer to Selecting Your Own Disks for FastResync.

VxVM Selecting the Disk for FastResync

(Refer to Selecting Your Own Disks for FastResync, if you want to select disks yourself.)

1. Select **Let Volume Manager decide which disks to use for FastResync**, then click **Next>** to bring up the **FastResync options** dialog.
2. You can now elect the options you require:

FastResync (DCO) Mirrors:	Provide the number of DCO volumes to be created.
Region Size (KB):	Size of region (in kilobytes).
Enable Dirty Region Logging	Check this box if you require DRL.

3. Click **Next>**. The Snapshot Type dialog appears.

Creating a Full Sized Snapshot

1. Select the type (**Full sized**). The Full Sized Options dialog appears.
2. Select **Create volume for snapshot** (refer to **Creating a Volume for Snapshot**) if you want to create a new volume; otherwise, select **Specify existing volume for snapshot** (refer to **Specifying an Existing Volume for Snapshot**).

Note The **Specify existing volume for snapshot** option will not be available (that is, it is grayed-out) if no existing volumes match the criteria for use as a snapshot (same size, started, no file system, not a current snap, instant ready, or if they are in the same disk group).

3. Check **Syncing** if you want the contents of the instant snapshot to be fully synchronized with the contents of the original volume at the point the snapshot is taken.

Creating Instant Snapshots

Creating a Space Optimized Snapshot

1. Select the type (**Space optimized**). Click **Next>**. The Space Optimized Options dialog appears.
2. Complete options for the cache object as follows:

Create a new cache object	Select this option if you want to create a new cache object.
Size	Define the size and select the unit for the new cache object.
Mirrors	Define the number of mirrors for the new cache object.
Autogrow	Select Autogrow if you require VxVM to automatically enlarge the cache if it is danger of overflowing. If Autogrow is enabled, but the cache cannot be extended; then VxVM disables the oldest and largest snapshot in the cache and releases the space to the cache.
Choose an existing cache object	Select this option if you want to use an existing cache object.
Cache objects	Select the existing cache object from the list provided.

When you have provided all necessary information in the dialog box, click **Next>**. The Snapshot options dialog appears.

3. Define a name for the snapshot, and add a comment if you want. Click **Finsh**.
4. The summary page is displayed.
 - Volume name
 - Name of new snapshot volume
 - Snapshot type (Space optimized)
 - Create cache (Yes or No)
 - Cache size
5. Click **Finish** if you are satisfied with your selections. Otherwise, click **<Back** to make changes.

Selecting Your Own Disks for FastResync

(Refer to VxVM Selecting the Disk for FastResync, if you want to VxVM to select the disks.)

1. Select **Manually select disks to use for FastResync** to bring up the **Select disks to use for FastResync** wizard.
2. Make your selections. The disks that you select should be in the right pane when you eventually click **Next>**.
3. You can now elect the options you require:

FastResync (DCO) Mirrors:	Provide the number of DCO volumes to be created.
Region Size (KB):	Size of region (in kilobytes).
Enable Dirty Region Logging	Check this box if you require DRL.

4. Click **Next>**. The Snapshot Type dialog appears.

Creating a Full Sized Snapshot

1. Select the type (**Full sized**). The Full Sized Options dialog appears.
2. Check **Syncing** if you want the contents of the instant snapshot to be fully synchronized with the contents of the original volume at the point the snapshot is taken.
3. Select **Create volume for snapshot** (refer to **Creating a Volume for Snapshot**) if you want to create a new volume; otherwise, select **Specify existing volume for snapshot Volume** (refer to **Specifying an Existing Volume for Snapshot**).

Note The **Specify existing volume for snapshot Volume** option will not be available (that is, it is grayed-out) if no existing volumes match the criteria for use as a snapshot (same size, started, no file system, not a current snap, instant ready, or if they are in the same disk group).

4. Click **Next>** when the confirmation window appears, and click **Finish** to create the instant snapshot.

Creating Instant Snapshots

Creating a Space Optimized Snapshot

1. Select the type (**Space optimized**). Click **Next>**. The Space Optimized Options dialog appears.
2. Complete options for the cache object as follows:

Create a new cache object	Select this option if you want to create a new cache object.
Size	Define the size and select the unit for the new cache object.
Mirrors	Define the number of mirrors for the new cache object.
Autogrow	Select Autogrow if you require VxVM to automatically enlarge the cache if it is danger of overflowing. If Autogrow is enabled, but the cache cannot be extended; then VxVM disables the oldest and largest snapshot in the cache and releases the space to the cache.
Choose an existing cache object	Select this option if you want to use an existing cache object. Note The option to specify an existing cache object is greyed-out if there are no existing caches.
Cache objects	Select the existing cache object from the list provided.

When you have provided all necessary information in the dialog box, click **Next>**. The Snapshot options dialog appears.

3. Define a name for the snapshot, and add a comment if you want. Click **Finsh**.
4. The summary page is displayed.
 - Volume name
 - Name of new snapshot volume
 - Snapshot type (Space optimized)
 - Create cache (Yes or No)
 - Cache size
5. Click **Finish** if you are satisfied with your selections. Otherwise, click **<Back** to make changes.

Creating a Full-sized Instant Snapshot of a Volume

▼ To create an instant full-sized snapshot of a volume

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Create**. The Create Instant Snapshot wizard appears.
3. Click **Next>**. The Snapshot Type dialog appears.
4. Select the type (**Full sized**). The Full Sized Options dialog appears.
5. Check **Syncing** if you want the contents of the instant snapshot to be fully synchronized with the contents of the original volume at the point the snapshot is taken.
6. Select **Create volume for snapshot** (refer to **Creating a Volume for Snapshot**) if you want to create a new volume; otherwise, select **Specify existing volume for snapshot Volume** (refer to **Specifying an Existing Volume for Snapshot**).

Note The **Specify existing volume for snapshot Volume** option will not be available (that is, it is grayed-out) if no existing volumes match the criteria for use as a snapshot (same size, started, no file system, not a current snap, instant ready, in the same disk group, and so on).

7. Click **Next>**.

Creating a Volume for Snapshot

In the **Choose the method by which to select Disks for this Volume** dialog, you can either let VxVM select the disks for you (**Let Volume Manager decide what disks to use for this volume**), refer to **VxVM Selecting the Disks**, or select your own disks (**Manually select disks to use for this volume**), refer to **Selecting Your Own Disks**.

VxVM Selecting the Disks

(Refer to **Selecting Your Own Disks**, if you want to select disks yourself.)

Creating Instant Snapshots

1. Select **Let Volume Manager decide which disks to use for this volume**, then click **Next>** to bring up the **Select the attributes for this volume** dialog.
2. Type an alphanumeric string in the **Volume Name** field. The total number of characters must be less than 32. This is a Volume Manager-specific name that is used in some Volume Manager commands. It is different from the volume label for the file system.
3. Select the volume type and attributes:
 - ◆ Concatenated (refer to Concatenated for a full description and corresponding attributes)
 - ◆ Striped (refer to Striped for a full description and corresponding attributes)
 - ◆ RAID-5 (refer to RAID-5 for a full description and corresponding attributes)
 - ◆ Mirrored (refer to Mirrored for a full description and corresponding attributes)If you choose Striped or RAID-5, Number of Columns and Stripe Unit Size need to have an entry. Defaults are provided.
4. Click **Next>** to continue. A summary of your selections appears. Click **<Back** to make changes; otherwise, click **Finish**.

Selecting Your Own Disks

(Refer to VxVM Selecting the Disks, if you want to VxVM to select the disks.)

1. Select **Manually select disks to use for this volume** to bring up the **Select disks to use for this volume** wizard. Then click **Next>** to bring up the **Select Disks to use with volume** dialog.
2. Make your selections. The disks that you select should be in the right pane when you eventually click **Next>**.
3. Specify the layout, select **Mirror Across** or **Stripe Across**; then select the device (Controller, Tray, Target, Enclosure). Select **Ordered** to use the specified storage to first concatenate the disks, then to form columns, and finally to form mirrors.
4. The **Select the attributes for this volume** screen appears.
5. Select the volume type and attributes:
 - ◆ Concatenated (refer to Concatenated for a full description and corresponding attributes)
 - ◆ Striped (refer to Striped for a full description and corresponding attributes)

- ◆ RAID-5 (refer to RAID-5 for a full description and corresponding attributes)
- ◆ Mirrored (refer to Mirrored for a full description and corresponding attributes)

If you choose Striped or RAID-5, Number of Columns and Stripe Unit Size need to have an entry. Defaults are provided.

6. Click **Next>** to continue. A summary of your selections appears:

Volume name
Snapshot type (Full sized)
Syncing (Yes or No)
Name of new snapshot volume
Disk selection
Volume type

7. Click **<Back** to make changes, otherwise, click **Finish**.

Specifying an Existing Volume for Snapshot

The existing volume must match the criteria (same size, started, no file system, not a current snap, instant ready, in the same disk group, and so on).

1. Select the volume to be used, then click **Next>**. A summary of the volume attributes appears.
2. Click **Finish** to accept the selection.

Creating Space-Optimized Snapshots of a Volume

▼ To create an instant space-optimized snapshot of a volume

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Create**. The Create Instant Snapshot wizard appears.
3. Click **Next>**. The Snapshot Type dialog appears.
4. Select the type (**Space optimized**).
5. Click **Next>**. The Space Optimized Options dialog appears.

Creating Instant Snapshots

6. Complete options for the cache object as follows:

Create a new cache object	Select this option if you want to create a new cache object.
Size	Define the size and select the unit for the new cache object.
Mirrors	Define the number of mirrors for the new cache object.
Autogrow	Select Autogrow if you require VxVM to automatically enlarge the cache if it is danger of overflowing. If Autogrow is enabled, but the cache cannot be extended; then VxVM disables the oldest and largest snapshot in the cache and releases the space to the cache.
Choose an existing cache object	Select this option if you want to use an existing cache object.
Cache objects	Select the existing cache object from the list provided.

When you have provided all necessary information in the dialog box, click **Next>**. The Snapshot options dialog appears.

7. Define a name for the snapshot, and add a comment if you want. Click **Finsh**.
8. The summary page is displayed.
 - Volume name
 - Name of new snapshot volume
 - Snapshot type (Space optimized)
 - Create cache (Yes or No)
 - Cache size
9. Click **Finish** if you are satisfied with your selections. Otherwise, click **<Back** to make changes.

Creating an Instant Snapshot of an ISP Volume

▼ To create an instant snapshot of an Intelligent Storage Provisioning (ISP) Volume

1. Select the ISP volume.
2. Choose **Actions > Instant Snapshot > Create**. The Create Instant Snapshot wizard appears.
3. Click **Next>**. The Snapshot Information window appears. The ISP Volume name and the ISP disk group name are listed.
 - ◆ Select a name for the snapshot and type it into the **Snapshot Name:** field.
 - ◆ Select the snapshot type: **Full-sized** or **Space Optimized**. Refer to [Creating Space-Optimized Snapshots of a Volume](#) or [Creating a Full-sized Instant Snapshot of a Volume](#).
 - ◆ Select the **Pool Name:** from the pool name pull-down menu.
4. Click **Next>**. If you chose a space optimized instant snapshot, complete options for the cache object as follows:

Create a new cache object	Select this option if you want to create a new cache object.
Size	Define the size and select the unit for the new cache object.
Mirrors	Define the number of mirrors for the new cache object.
Autogrow	Select Autogrow if you require VxVM to automatically enlarge the cache if it is danger of overflowing. If Autogrow is enabled, but the cache cannot be extended; then VxVM disables the oldest and largest snapshot in the cache and releases the space to the cache.
Choose an existing cache object	Select this option if you want to use an existing cache object.
Cache objects	Select the existing cache object from the list provided.

5. When you have provided all necessary information in the dialog box, click **Next>**. The Select User Template dialog box appears.

Creating Instant Snapshots

6. You are now asked to select an existing user template, or if you want to define a new template (Create the User Template).
7. Select the capabilities you want. (Refer to Select Capabilities, Using Intelligent Storage Provisioning.) Click **Next**>.
8. The summary page is displayed.
 - Volume name
 - Name of new snapshot volume set
 - Snapshot type (Space optimized)
 - Create cache (Yes or No)
 - Cache size
9. Click **Finish** if you are satisfied with your selections. Otherwise, click <**Back** to make changes.

Creating an Instant Snapshot of an ISP Volume Set

▼ To create an instant snapshot of an Intelligent Storage Provisioning (ISP) Volume Set

1. Select the ISP volume set.
2. Choose **Actions > Instant Snapshot > Create**. The Volume Set Snapshot wizard appears.
3. Click **Next**>. The Volume Set Snapshot dialog appears. You can choose to create a New Volume Set Snapshot (space optimized), or Create a Snapshot using an existing volume set.
4. Click **Next**>. If you chose a space optimized instant snapshot, complete options for the cache object as follows:

Choose an existing cache object	Select this option if you want to use an existing cache object.
Cache objects	Select the existing cache object from the list provided.
Create a new cache object	Select this option if you want to create a new cache object.

Size	Define the size and select the unit for the new cache object.
High Water Mark	When cache usage reaches this value, and the new required cache size would exceed the value of Max Grow, the oldest snapshot in the cache is deleted. If there are several snapshots with the same age, the largest of these is deleted.
Max Grow	Limits the maximum size to which a cache can grow. To estimate this size, consider how much the contents of each source volume are likely to change between snapshot refreshes, and allow some additional space for contingency.
Add by	Value by which the cache is increased in size.
Autogrow	Select Autogrow if you require VxVM to automatically enlarge the cache if it is danger of overflowing.

5. When you have provided all necessary information in the dialog box, click **Next>**. The Select User Template dialog box appears.
6. You are now asked to select an existing user template, or if you want to define a new template (Create the User Template).
7. When you have provided all necessary information in the dialog box, click **Next>**. The Select Capabilities dialog box appears.
8. When you have selected all of the capabilities you need for this volume set, click **Next>**. The Set Configuration Rules dialog appears.
9. When you have selected all of the configuration rules you need for this volume set, click **Next>**. The Volume Set Information dialog appears.
10. Enter the name of the volume set of which you are creating a snapshot, add comments and names to all volumes within the volume set, and then click **Next>**.
11. The summary page is displayed.
 - Volume set name
 - Disk group name
 - Storage pool name
 - Snapshot Volume Set name
 - Selected Capabilities
 - Selected Rules

Creating Instant Snapshots

12. Click **Finish** if you are satisfied with your selections. Otherwise, click **<Back** to make changes.

Creating a Full-sized Instant Snapshot of a Volume Set

▼ To create an instant full-sized snapshot of a volume set

1. Select the volume set.
2. Choose **Actions > Instant Snapshot > Create**. The Create Instant Snapshot wizard appears.
3. Click **Next>**. The Snapshot Type dialog appears.
4. Select the type (**Full sized**). The Full Sized Options dialog appears.
5. Select the **volume sets** you where you want the instant snapshot to reside.
6. Check **Syncing** if you want the contents of the instant snapshot to be fully synchronized with the contents of the original volume at the point the snapshot is taken.
7. Click **Next>**.

Creating Space-Optimized Snapshots of a Volume Set

▼ To create an instant space-optimized snapshot of a volume set

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Create**. The Create Instant Snapshot wizard appears.
3. Click **Next>**. The Snapshot Type dialog appears.
4. Select the type (**Space optimized**). Click **Next>**. The Space Optimized Options dialog appears.

5. Complete options for the cache object as follows:

Create a new cache object	Select this option if you want to create a new cache object.
Size	Define the size and select the unit for the new cache object.
Mirrors	Define the number of mirrors for the new cache object.
Autogrow	Select Autogrow if you require VxVM to automatically enlarge the cache if it is danger of overflowing. If Autogrow is enabled, but the cache cannot be extended; then VxVM disables the oldest and largest snapshot in the cache and releases the space to the cache.
Choose an existing cache object	Select this option if you want to use an existing cache object.
Cache objects	Select the existing cache object from the list provided.

When you have provided all necessary information in the dialog box, click **Next>**. The Snapshot options dialog appears.

6. Define a name for the snapshot, and add a comment if you want. Click **Finsh**.
7. The summary page is displayed.
 - Volume name
 - Name of new snapshot volume set
 - Snapshot type (Space optimized)
 - Create cache (Yes or No)
 - Cache size
8. Click **Finish** if you are satisfied with your selections. Otherwise, click **<Back** to make changes.

Creating Instant Snapshots

Notes:

- ◆ A Volume Set must be identical in size and layout as well as enabled for FastResync in order to be used as a new volume set snapshot.
- ◆ Volume sets for breaking off snapshots must have the same limitations as volume set full-sized snapshots.
- ◆ You cannot add or delete volumes in a volume set snapshot and you cannot add snapshot volumes to a volume set.

Creating Break-Off Snapshots

You can create traditional third-mirror snapshots.

▼ To create a third-mirror snapshot

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Create**. The Create Instant Snapshot wizard appears.
3. Click **Next>**. The Snapshot Type dialog appears.
4. Select the type (**Break off**). Click **Next>**. The Full Sized Break Off Options dialog appears.
5. Complete the Full Sized Break Off Options dialog as follows:

Specify number of mirrors:	Number of mirrors to be broken off.
Specify mirror objects	Select the mirrors from the available mirrors. Note This option is greyed-out unless the plex state is SNAPDONE .
Selected Mirrors	Move the mirrors to be added from Available mirrors to Selected mirrors.

6. When you have provided all necessary information in the dialog box, click **Next>**.

7. The summary page is displayed. Click **Finish** if you are satisfied with your selections. Otherwise, click **<Back** to make changes.

Disabling FastResync

Use the following procedure to disable FastResync:

▼ To disable FastResync for a volume

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Disable FastResync**.
3. When the confirmation dialog appears (Are you sure you want to disable FastResync for *volume-name*?), click **Yes** to confirm.
4. FastResync is then disabled. This removes the Disk Change Object (DCO) and DRLs, if any.

Note This option is disabled if the volume is in a snapshot chain.

Refreshing an Instant Snapshot

You can immediately retake an instant snapshot or instant space-optimized snapshot at any time by using the **Refresh** procedure. If a fully synchronized instant snapshot is required, you must wait for the new resynchronization to complete.

Refreshing an instant snapshot replaces it with another point-in-time copy of a parent volume.

If the source volume is not specified, the immediate parent of the snapshot is used.

Note The snapshot being refreshed must not be open to any application. For example, any file system configured on the volume must first be unmounted.

▼ To refresh a volume

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Refresh**
3. Select **Yes** from the Refresh confirmation message.

Reattaching an Instant Snapshot

Note This operation is not supported for instant space-optimized snapshots.

Use the **Reattach** procedure to reattach the plexes of an instant snapshot volume to the specified volume in a snapshot hierarchy, and synchronizes them from that volume.

▼ To reattach a volume

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Reattach**. The Snapshot Reattach dialog appears.
3. Specify the volume to which the snapshot is to be reattached.
4. Specify the number of mirrors to reattach.
5. Select **OK** to confirm the reattachment.

Note The snapshot being reattached must not be open to any application. For example, any file system configured on the snapshot volume must first be unmounted.

Restoring a Volume from an Instant Snapshot

On occasion, it may be desirable to reinstate the contents of a volume from a backup or modified replica within a snapshot volume. You can use the **Restore** procedure to achieve this.

For a full instant snapshot, some or all of its plexes may be reattached to the parent volume or to a specified source volume in the snapshot hierarchy above the snapshot volume. If the Destroy checkbox is checked, all the plexes of the snapshot are reattached and the snapshot volume is removed.

For an instant space-optimized snapshot, the cached data is used to restore the contents of the specified volume. The snapshot itself remains unchanged by the operation.

Note For this operation to succeed, the volume that is being restored and the snapshot volume must not be open to any application. For example, any file systems that are configured on either volume must first be unmounted.

▼ To restore a volume

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Restore**. The Snapshot Restore dialog appears.
3. Specify the snapshot to be restored.
4. Check the **Synchronize** box to initiate volume synchronization.
5. Attach the snapshot plexes to the parent upon completion of the volume synchronization.
6. Check **Destroy** if the snapshot volume is to be removed after reattaching all the plexes.
7. Select **OK** to confirm the restore.

Dissociating an Instant Snapshot

Note This operation is not supported for instant space-optimized snapshots.

Use the **Dissociate** procedure to dissociate an instant snapshot from its parent volume or volume set in the snapshot hierarchy. If the snapshot has a dependent snapshot hierarchy below it, this hierarchy is adopted by the parent volume.

Note The snapshot volume or volume set and any snapshots in the hierarchy must have been fully synchronized for this procedure to succeed.

▼ To dissociate a snapshot from a volume

1. Select the snapshot.
2. Choose **Actions > Instant Snapshot > Dissociate**. The Snapshot Confirm Dissociate dialog appears.
3. Select **Yes** to confirm the dissociate.

▼ To dissociate a snapshot from a volume set

Note Dissociating a volume set can cause inconsistencies in the snapshot hierarchy in the unlikely event of a failure.

Creating Instant Snapshots

1. Select the snapshot from the volume set.
2. Choose **Actions > Instant Snapshot > Dissociate**.
3. When the Snapshot Confirm Dissociate dialog appears, (Are you sure you want to dissociate THIS_VSET)? Click **Yes** to confirm.

Splitting an Instant Snapshot Hierarchy

Note This operation is not supported for instant space-optimized snapshots.

The topmost snapshot volume in the hierarchy must have been fully synchronized for this task to succeed. Snapshots that are lower down in the hierarchy need not have been fully resynchronized.

Use the **Split** procedure to split the snapshot hierarchy into two independent hierarchies.

▼ To split a snapshot hierarchy

1. Select the snapshot.
2. Choose **Actions > Instant Snapshot > Split**. The Snapshot Confirm Split dialog appears.
3. Select **Yes** to confirm the split.

▼ To split a snapshot hierarchy of a volume set

1. Select the snapshot of the volume set.
2. Choose **Actions > Instant Snapshot > Split**.
3. When the Snapshot Confirm Split appears, (Are you sure you want to split THIS_VSET?) click **Yes** to confirm.

Note Splitting a volume set can cause inconsistencies in the snapshot hierarchy in the unlikely event of a failure.

Start Synchronization

The **Start Synchronization** procedure synchronizes the contents of an instant snapshot with the contents of the original volume at the point in time that the snapshot was taken. If you wish to move an instant snapshot volume to another disk group for export to another machine for off-host processing, you must ensure that the snapshot volume has been completely synchronized.

If synchronization of an instant snapshot was not started when the snapshot was created, or synchronization was paused, you can start synchronization.

Note Synchronization of the contents of a snapshot with its original volume is not possible for instant space-optimized snapshots.

Use the **Start Synchronization** procedure to start synchronization.

▼ To start synchronization

1. Select the snapshot.
2. Choose **Actions > Instant Snapshot > Start Synchronization**. The Start Synchronization Confirm dialog appears.
3. Select **Yes** to confirm synchronization.

Note You can pause, restart, or abort synchronization from the Task View. (Select the task, then right click).

Adding Maps

You can add a map to a volume that has already been prepared for instant snapshot.

When you add a map to volume, you can either select your own disks, or let VxVM select the disks for you:

- ◆ Selecting Your Own Disks
- ◆ VxVM Selecting the Disks

Selecting Your Own Disks

(Refer to VxVM Selecting the Disks, if you want to VxVM to select the disks.)

Use the following procedure to add a map to a volume, selecting your own disks:

Using the Full Volume Snapshot Feature

▼ To add a map to a volume

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Add map**. The Add map dialog appears.
3. Choose the number of maps to add.
4. Choose **Manually select disks**. Then move the required disks from **Available disks:** list to the **Selected disks:** list.
5. Click **OK** to complete the operation.

VxVM Selecting the Disks

(Refer to *Selecting Your Own Disks*, if you want to select the disks yourself.)

Use the following procedure to add a map to a volume, letting VxVM select the disks.:

▼ To add a map to a volume

1. Select the volume.
2. Choose **Actions > Instant Snapshot > Add map**. The Add map dialog appears.
3. Choose the number of maps to add.
4. Choose **Auto select disks**. Click **OK** to complete the operation.

Using the Full Volume Snapshot Feature

The following five snapshot procedures are supported by VEA for creating snapshots of pre-4.0 (pre-110) disk group volumes. These procedures are accessed through the Volume context menu:

- ◆ **Snap Start - Creating a Snapshot Copy of a Volume**
Attaches the snapshot mirror to the original volume.
- ◆ **Snap Shot - Creating the New Volume**
Detaches the snapshot mirror from the original volume and creates a snapshot volume that is associated with the snapshot mirror.
- ◆ **Snap Back - Merging a Snapshot Volume (Volume Snapback)**

Reattaches the snapshot mirror to the original volume.

- ◆ **Snap Clear - Dissociating a Snapshot Volume (Volume Snap Clear)**

Permanently removes the snapshot mirror from the original volume and releases its space.

- ◆ **Snap Abort - Aborting a Snapshot**

Aborts the snapshot operation after issuing the **Snap Start** command. **Snap Abort** permanently removes the snapshot mirror from the volume and releases its space.

The five snapshot commands are subcommands of the **Snap** command on the Volume context menu.

(Refer to Using the Instant Snapshot Feature for information on creating snapshots of 4.0 (version 110) disk group volumes.)

Using Snapshot

In summary, the snapshot procedure consists of three steps:

1. Attaching a snapshot mirror to a volume, using **Snap Start**.
2. Detaching the snapshot mirror and creating the snapshot volume that is associated with the snapshot mirror. The **Snap Shot** command is used for this step.
3. Either reattaching the snapshot to the volume after you are done using it, or permanently disconnecting it from the volume. **Snap Back** reattaches the snapshot mirror; **Snap Clear** permanently disconnects the snapshot mirror from the volume, removing the association between the snapshot volume and the original volume.

Creating a Snapshot Copy of a Volume

Snap Start creates a snap mirror or plex, which is attached to and synchronized with a volume. Once the snap mirror is synchronized, its status, as displayed on the Mirrors tab of the GUI, changes to Snap Ready.

After the snap mirror is synchronized with the volume, it can be used as a snapshot mirror. The synchronization process can take a while, but it does not interfere with use of the volume. If the Snap Start process fails, the snapshot mirror is removed and its space is released. After the mirror is synchronized, it continues to be updated until it is detached.

Note To create a snapshot mirror on a volume, you must have an empty disk in the disk group of that volume.

Using the Full Volume Snapshot Feature

▼ To create a snapshot mirror, do the following:

1. Right-click on the volume that you want to take a snapshot of.
A context menu is displayed.
2. Select **Snap > Snap Start**.

A new snapshot mirror is attached to the volume.

The new mirror is added to the Mirrors tab for the volume. The mirror is identified as a Snapshot plex and is identified with the Snapshot icon. After the mirror is synchronized with the volume, its status becomes Snap Ready.

The DCO volume is created to track the regions on a volume that are changed while a mirror is detached. The DCO volume is not included in the tree view of the Volume Manager GUI. To view the DCO volume, you must use the Disk View. To access the Disk View, click the **Disk View** tab in the right pane or select Disk View from a disk's context menu.

Note The **Break Mirror** and **Remove Mirror** features do not work with the snapshot mirror.

Creating the New Volume

Snap Shot detaches the snapshot mirror and creates a new volume that is associated with that mirror. This process takes a relatively short time, typically less than a minute, during which users can be requested to not access the original volume. After the snapshot volume is detached, it can be used for backup or other purposes.

▼ To take a snapshot, do the following:

1. Right-click on the volume with the snapshot mirror.
A context menu is displayed.
2. Select **Snap > Snap Shot**.

The Snapshot Volume dialog box is displayed. You are asked to select the snapshot mirror to be used to create the snapshot volume.

3. Select the mirror and click **OK**.

When there is only one snapshot mirror, "Automatic select mirror to add snapshot" will be highlighted. If there are multiple mirrors and you want to select the mirror to be used, click to select "Manual select which mirror to add snapshot." You then can select the mirror from the list on the screen.

Click **OK**.

The snapshot mirror is detached from the original volume and a new volume is created that is associated with the snapshot mirror.

The snapshot mirror is no longer displayed on the Mirrors tab for the original volume. The new snapshot volume is displayed under the Volumes folder in the tree view. The program assigns it the next available drive letter. You can now use the snapshot volume for backup or other purposes.

Notes:

- ◆ FastResync is not supported for RAID-5 volumes
- ◆ The **Snapshot Volume** dialog requires a volume name and a snapshot name.
- ◆ The snapshot volume is an exact copy of the original volume at the time the snapshot phase of the procedure starts.
- ◆ After you create the snapshot copy of the volume, remember to back it up to tape or some other media. Remove the snapshot volume when it is no longer needed.
- ◆ The snapstart phase of this task may take a long time, depending on the size of the volume.
- ◆ If you remove the snapshot mirror, you cannot create a snapshot volume until you repeat the snapstart phase of the procedure.

Merging a Snapshot Volume (Volume Snapback)

The volume **Snap Back** procedure merges a snapshot copy of a volume with the original volume. The snapshot plex is detached from the snapshot volume and attached to the original volume. The snapshot volume is removed. This task resynchronizes the data in the volume so that the plexes are consistent. By default, the data in the original plex is used for the merged volume.

▼ **To reattach a snapshot mirror to the original volume, do the following:**

1. Right-click the volume whose snapshot you want to reattach.
A context menu is displayed.
2. Select **Snap > Snap Back**.
3. Specify whether the snapshot volume is resynchronized with the original volume, or the original volume is synchronized to the snapshot volume and click **OK**.

Using the Full Volume Snapshot Feature

Note One situation where you might want to resynchronize to the snapshot volume rather than the original volume is the case where something has happened to the original volume so that its data integrity is no longer sound.

The snapshot mirror is reattached to the original volume and the snapshot volume is deleted if it is the last plex of the volume. (If the snapshot volume has other mirrors associated with it, it is not deleted.)

Dissociating a Snapshot Volume (Volume Snap Clear)

The volume **Snap Clear** procedure permanently breaks the association between a snapshot and its original volume. The snapshot volume becomes an independent volume.

Snap Clear permanently removes a snapshot mirror from a volume. Snap Clear can be specified only after the snapshot command has been issued.

▼ **To permanently remove a snapshot mirror from a volume, do the following:**

1. Right-click on the volume, whose snapshot mirror you are removing.

A context mirror is displayed.

2. Select **Snap > Snap Clear**.

A message box is displayed asking you to confirm the snap clear operation for the specified volume.

3. Click **Yes**.

The snapshot mirror is permanently disconnected from the original volume and is no longer displayed on the Volume Manager console.

Aborting a Snapshot

The volume **Snap Abort** procedure allows you to abort the snapshot procedure before it completes. If you have started the procedure, through **Snap Start**, but have not yet clicked on **Snap Shot** to complete the procedure, you can select **Snap Abort** to abort the procedure.

▼ **To abort a snapshot operation begun with snap start, do the following:**

1. Right-click on the volume that you want to abort the snapshot operation on.

A context menu is displayed.

2. Select **Snap > Snap Abort**.

A message box is displayed asking you to confirm the snap abort operation for the specified volume.

3. Click OK.

The snapshot process is aborted.

Displaying the Snapshots for a Volume

You can list the direct snapshots belonging to a volume

▼ Listing the Snapshots for a Volume

1. Left-click on the volume. The volume details are displayed.
2. If the volume has direct snapshots, the **Snapshot Volumes** tab is displayed.
3. Click on the **Shapshot Volumes** tab.
4. The snapshots belonging to the volume are displayed in the volume table.

Displaying Cache Volume Information

You can display the backing cache volume information for space optimized volumes.

▼ Listing Backing Cache Information

1. Left-click on the volume. The volume details are displayed.
2. If the volume is space optimized, the **Cache Objects** tab is displayed.
3. Click on the **Cache Objects** tab.
4. Cache information for the volume is displayed in the volume table.

Backing up a Volume

To back up a volume, use the following procedure:

▼ Backing up a volume using a space-optimized snapshot

1. If you intend to take an instant space-optimized snapshot of the volume, you may wish to consider first setting up a shared cache object in the same disk group as the volume. If you intend to split the volume and snapshot into separate disk groups (for example, to perform off-host processing), you must use a fully synchronized full-sized snapshot (which does not require a cache object). You cannot use a space-optimized snapshot for this purpose.
2. If you have not already done so, enable FastResync on the volume. Click:

Actions > Instant Snapshot > Enable FastResync

The **Create** task creates an instant snapshot that is immediately available for making a backup. Read requests for data in the original volume at the point in time that the snapshot was taken are satisfied by reading either from a non-updated region of the original volume, or from the copy of the original contents of an updated region that have been recorded by the snapshot.

Note Unless you perform the **Start Synchronization** procedure, the contents of an instant snapshot are unlikely ever to become fully synchronized with the contents of the original volume at the point in time that the snapshot was taken. If you wish to move an instant snapshot volume to another disk group for export to another machine for off-host processing, you must ensure that the snapshot volume has been completely synchronized.

For instant space-optimized snapshots that share a cache object, the specified region size must be greater than or equal to the region size specified for the cache object.

3. Use the **Create** task create an instant snapshot that is immediately available for making a backup. If you specify a name and size for a storage cache, a space-optimized snapshot is created that requires less space than does a full-sized snapshot volume.

Note By default, VEA attempts to avoid placing snapshot plexes on a disk that already holds any plexes of a data volume. However, this may be impossible if insufficient space is available in the disk group. In this case, VEA uses any available space on other disks in the disk group. If the snapshot plexes are placed on disks which are used to hold the plexes of other volumes, this may cause problems when you subsequently attempt to move a snapshot volume into another disk group.

You have the following choices for what to do with an instant snapshot:

- ◆ Refresh the contents of the snapshot. This effectively creates a new snapshot of the original volume ready for another backup. If synchronization was already in progress on the snapshot, this operation may result in large portions of the snapshot having to be resynchronized.
- ◆ Restore the contents of the original volume from the snapshot volume. An instant space-optimized snapshot always remains intact at the end of the operation.

▼ Backing up a volume using a full-sized snapshot


1. For full-sized instant snapshots, you must create a volume that is to be used as the snapshot volume. This must be the same size as the volume for which the snapshot is being created, and it must also have the same region size. Alternatively, if there is a sufficient number of suitable plexes available in the volume, you can break these off and use them to create a full-sized instant snapshot.

Note To turn one or more existing plexes in a volume into a break-off instant snapshot volume, the volume must be a non-layered volume with a `mirror` or `mirror-stripe` layout, or a RAID-5 volume that you have converted to a special layered volume. The volume must have a sufficiently number of available plexes that are in the `ACTIVE` state. The plexes in a volume with a `stripe-mirror` layout are mirrored at the sub-volume level, and cannot be broken off.



2. Reattach some or all of the plexes of the snapshot volume with the original volume.
3. Restore the contents of the original volume from the snapshot volume. You can choose whether none, a subset, or all of the plexes of the snapshot volume are returned to the original volume as a result of the operation.
4. Dissociate the snapshot volume entirely from the original volume. This may be useful if you want to use the copy for other purposes such as testing or report generation. If desired, you can delete the dissociated volume.
5. If the snapshot is part of a snapshot hierarchy, you can also choose to dissociate this hierarchy from its parent volumes.

Preparing to Restore a Volume From Backup

The **Volume Prepare Restore** procedure stops the volume, sets the volume to an uninitialized state, and restarts the volume (without resynchronizing the volume mirrors). The volume can then be reloaded from backup (refer also to *Merging a Snapshot Volume (Volume Snapback)*).



Preparing to Restore a Volume From Backup



File System Tasks

7

File system tasks include:

- ◆ Creating a New File System on a Volume
- ◆ Creating a New File System on a Volume Set
- ◆ Mounting a File System
- ◆ Unmounting a File System
- ◆ Removing a File System from the File System Table
- ◆ Defragmenting a File System
- ◆ Creating a Snapshot Copy of a File System
- ◆ Removing a Snapshot Copy of a File System
- ◆ Checking a File System
- ◆ Viewing File System Properties
- ◆ Using the QuickLog Feature
- ◆ Unmounting a File System from a Cluster Node
- ◆ Removing Resource Information for a Cluster File System
- ◆ Maintaining Storage Checkpoints
- ◆ Using Multi-Component File Systems

File System Task Roadmap

The following table provides a summary of common VEA file system tasks.

To perform this task ...	Use this procedure ...
<i>File System Setup Tasks:</i>	
Add a new file system to a volume	Creating a New File System on a Volume
Add a new file system to a volume set	Creating a New File System on a Volume Set
Mount a file system	Mounting a File System
<i>Maintenance Tasks:</i>	
Defragment a file system	Defragmenting a File System
Create a snapshot copy of a file system for backup	Creating a Snapshot Copy of a File System
Remove a snapshot copy of a file system	Removing a Snapshot Copy of a File System
Maintaining Storage Checkpoints	Maintaining Storage Checkpoints
<i>Recovery Tasks:</i>	
Perform a file system consistency check	Checking a File System
<i>Other Tasks:</i>	
Unmount a file system	Unmounting a File System
Cluster mounting a file system	Mounting a File System
Removing cluster file system resource information	Removing Resource Information for a Cluster File System
Create a QuickLog volume	Creating a QuickLog Volume
Enable a QuickLog device	Enabling a QuickLog Device for a File System
Disable a QuickLog device	Disabling a QuickLog Device for a File System
Remove a QuickLog volume	Removing a QuickLog Volume

Creating a New File System on a Volume

You can use the Volume Manager **New File System** procedure to place a file system onto an existing volume and mount the file system. This also updates the file system table file, if necessary.

▼ To add a new file system to an existing volume

1. Select the volume to contain the file system.
2. Choose **Actions > File System > New File System**.
3. Complete the New File System dialog box as follows:

File system type:	Select the file system type from the options shown.
Create Options:	<ul style="list-style-type: none"> ◆ See Allocation unit/Block size for allocation and block and settings. ◆ To specify other create details, click New File System Details and specify any additional options in the New File System Details dialog box.
Compress (checkbox)	<ul style="list-style-type: none"> ◆ If your platform supports file compression, then you can select Compress to compress the files on your file system.
Mount Options:	<ul style="list-style-type: none"> ◆ Enter the mount point for the file system, if you want the file system mounted at system startup. ◆ Select the Create mount point checkbox if you want the system to create the mount point if it does not already exist. ◆ Select the Read only and Honor setuid checkboxes, as required. ◆ Select the Add to file system table and Mount at boot checkboxes to update the system table file and mount the file system at system startup. ◆ Select the Add to file system table checkbox, and leave the Mount at boot checkbox unselected to update the system table file, and not mount the file system at system startup. ◆ To specify mount details, click Mount File System Details and specify any additional options using the checkboxes and Extra Options entry in the Mount Details dialog box.

After you have provided all necessary information in the dialog box, click **OK**.

Creating a New File System on a Volume

Notes:

- ◆ The **New File System** task requires a file system type.
- ◆ If the file system is to be mounted at startup, a file system mount point is required.
- ◆ If the Add to file system table checkbox is selected, then the file system table is automatically updated when the file system is mounted.
- ◆ The mount point must be an absolute pathname (that is, it must begin with /).
- ◆ In VxFS 4.0 default disk layout is version 6. However older layouts 4 and 5 are supported.

Allocation unit/Block size

File System	Allocation Unit Values	Block Size Values
ufs	1024, 2048, 4096, 8192	4096, 8192
vxfs	same as block size	1024, 2048, 4096, 8192

New File System Details

Option	File Systems	Description
Checkbox - Support large file size (Greater than 2 GB)	vxfs, jfs (Journalled File System)	Check if file system contains files greater than 2 GBytes.
Checkbox - Specify activity logging area size	vxfs	Check if you want to specify the activity logging area size. If checked, you need to provide a value for the log size (value must be expressed as the number of blocks, in the range 256 to 262144, inclusive).
Extra options	vxfs	See mkfs(1) and mount(1) manual pages for option details.

Mount File System Details

The Mount File System Details dialog box contains the following sections:

- ◆ Disable Quick I/O checkbox - available only with vxfs file systems. This disables the VERITAS Quick I/O for Databases option.
- ◆ Support large file sizes checkbox - available only with vxfs file systems. Check this box if the file system contains files greater than 2 GBytes.
- ◆ Quicklog device checkbox - available only with vxfs file systems. Check this box to enable the file system to use VERITAS QuickLog. Select also the log device number to be used (1 to 31).

(QuickLog is available only for layouts 4 and 5. Use Intent Log for layout 6.)

- ◆ Cluster Options
- ◆ Permissions/Ownership
- ◆ Advanced Options

Cluster Options

Cluster Options Area

Option	File Systems	Description
Checkbox - Cluster mount. This is present only if the optional cluster functionality is available.	vxfs	Mounts the file system in shared mode. Select the nodes to be mounted.
Once the Cluster mount checkbox is selected, you can set the Node mount information, primary node, and Service group information.	vxfs	Mounts the file system in shared mode. Select the nodes to be mounted, and their mode (Read Only). Select the Service group (cfs, by default). This defines the vcs service group that the mount resources will go into. You can select an existing group, or type in a new one.

Note If you do not specify a service group, a new service group is created. The name of the generated service group takes the form `vrts_vea_cfs_xx` where `xx` is a running serial number. If you specify a service group which does not exist, a new service group is created. If the service group does exist, it checks the node lists specified in the command with the existing node lists in the group. If the node lists do not match the service group selection operation will fail.

Creating a New File System on a Volume

Permissions/Ownership

Permissions/Ownership Area

Option	File Systems	Description
Permissions/Ownership (applies only to local file systems)	vxfs	Note To set ownership, both group and user must be specified; and the user must belong in the group.

Advanced Options

Advanced Options Area

Option	File Systems	Description
log/delaylog/tmplog	jfs, jfs2	Controls intent logging. To maintain file system integrity after a system failure, logging must be enabled.
Checkbox - Clear all data extents before file allocation.	vxfs	Clears all data extents before file allocation.
Checkbox - File system caching policy	vxfs	Note that, while the tmpcache option to mincache improves performance, it reduces integrity.
Checkbox - Disable access time recording	vxfs	Disables access time recording.

Mounting a File System

The **Mount File System** procedure mounts a file system and updates the file system table file, if necessary.

▼ To mount a file system on an existing volume

1. Select the volume that contains the file system to be mounted.
2. Choose **Actions > File System > Mount File System**.
3. Complete the Mount File System dialog box

▼ To mount any file system

1. Select the file system to be mounted.
2. Choose **Actions > Mount File System**.
3. Complete the Mount File System dialog box as follows:

FS Type:	Select the file system type.
Mount Options: (refer to Creating a New File System on a Volume for available options)	<ul style="list-style-type: none"> ◆ If you want the system to use the mount options defined in the system table, check Mount using options in the file system table. ◆ Enter the mount point for the file system, if you want the file system mounted at system startup. ◆ Select the Create mount point checkbox if you want the system to create the mount point if it does not already exist. ◆ Select the Read only and Honor setuid checkboxes, as required. ◆ To specify mount details, click Mount File System Details and specify any additional options using the checkboxes and Extra Options entry in the Mount Details dialog box (see Advanced Options Area).

After you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The mount point must be an absolute pathname (that is, it must begin with /).

Creating a New File System on a Volume Set

You can use the Volume Manager **New File System** procedure to place a file system onto an existing volume set and mount the file system. This also updates the file system table file, if necessary.

▼ **To add a new file system to an existing volume set**

1. Select the volume set to contain the file system.
2. Choose **Actions > File System > New File System**.
3. Proceed as for creating a new file system on a volume (Creating a New File System on a Volume).

Note The volume set device will be shown as a child of the file system instance in the tree.

Unmounting a File System

The **Unmounting File System** procedure unmounts a file system.

▼ **To unmount a file system on a volume**

1. Select the volume containing the file system to be unmounted.
2. Choose **Actions > File System > Unmount File System**.
3. Click **Yes** in the Unmount File System dialog to confirm that you want the file system unmounted.
4. If an entry exists for this file system in the file system table, a Remove File System dialog appears. Click **Yes** in the Remove File System dialog if you want the file system table entry removed.

Removing a File System from the File System Table

The file systems **Remove from File System Table** procedure removes its resource information from the file system table.

▼ To remove file system information from the file system table

1. Select the file system to be removed.
2. Choose **Actions > Remove from File System Table**.
3. Click **Yes** in the Remove File System dialog to confirm that you want the file system information removed.

Defragmenting a File System

The **Defrag File System** procedure defragments a file system and the underlying volume, removing disjointed extents, and so on.

▼ To defragment a file system on a volume

1. Select the volume containing the file system to be defragmented.
2. Choose **Actions > File System > Defrag File System**.
3. Select **Yes** in the displayed dialog box

▼ To defragment any file system

1. Select the file system to be defragmented.
2. Choose **Actions > Defrag File System**.
3. Select **Yes** in the displayed dialog box.

Creating a Snapshot Copy of a File System

The **File System Snapshot** procedure prepares VxFS for online backup by creating a snapshot copy of the file system. This is done by creating a new volume, mounting a snapshot file system on the new volume, and copying the contents of the original file system to the snapshot file system. The snapshot file system is a consistent copy of the original file system at the time the snapshot was made. You can use the snapshot file system to make a backup of the original file system at a convenient time. After the backup is made, you can remove the snapshot file system by removing the underlying snapshot volume.

▼ To create a snapshot copy of file system data

1. Select the file system to be backed up.
2. Choose **Actions > Snapshot > Create...**
3. Complete the Snapshot File System dialog box as follows:

New Volume Name:	By default, the new volume name displayed is composed of the name of the file system to be backed up and the suffix <code>_vxfsnap</code> (for example, Vol5_vxfsnap). You can replace the default name with the volume name of your choice.
Snapshot Mount Point:	By default, the new snapshot mount point displayed is composed of the mount point of the file system to be backed up and the suffix <code>_vxfsnap</code> (for example, <code>/Vol5_vxfsnap</code>). You can replace the default name with the mount point of your choice.
Snapshot Size:	Enter the size of the snapshot copy of the file system.
Cluster node (displayed only if your node is part of a cluster):	If your node is part of a cluster, you can now select the node in the cluster to contain the file system. If the node you select is not displayed in the current window, then it resides on a remote node. In order to remove this snapshot, you will have to open the remote node and display the created file system before removing it.
Select disks to use for the volume:	Select the disks for Volume screen. The default setting is for Volume Manager to assign the disks for you. To manually select the disks, click the "Manually select which disks to use for volume" radio button. The disks that you select should be in the right pane when you click Apply or let VxVM decide which disks to use for the volume.

After you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ Snapshots are shown in the left pane tree view as children of the file system they are based on. They are shown in the right pane of the table view of the File Systems at the same level as the file system they are based on.
- ◆ In the case of creating a snapshot of a cluster file system, snapshot creation can be run only on the master node
- ◆ Snapshots cannot be made of Multi-component file systems.
- ◆ The **File System Snapshot** task is only available for VxFS file systems.
- ◆ The snapshot file system is an exact copy of the original file system at the time the snapshot was created.
- ◆ After you create the snapshot copy of the file system, remember to back it up to tape or some other media. Remove the snapshot file system when it is no longer needed.
- ◆ Users can continue to use the original file system during this snapshot backup procedure.
- ◆ Refer to the *VERITAS File System Administrator's Guide* for guidelines on choosing the snapshot file system size.

Removing a Snapshot Copy of a File System

The removal of a snapshot file system is performed in two steps:

1. The system first unmounts the snapshot file system and, in the case of a cluster file system, deletes the snapshot mount point from the cluster configuration.
2. The system then deletes the volume that was created for the snapshot file system. This step must be executed on the master node in the case of a cluster file system.

Use the following procedure to remove a snapshot copy of a file system.

▼ **To remove a snapshot copy of file system data**

1. Select the snapshot file system to be removed.
2. Choose **Actions > Snapshot > Remove**
3. Click **Yes** in the Remove Snapshot File System dialog to confirm that you want the snapshot deleted.

Checking a File System

Note If you are removing the snapshot on a slave node. The message “Error V-40-49408-40 VxVM exedit ERROR V-5-1-0 operation must be executed on master” may appear. In this case, you need to go to the master node and delete the snapshot volume after the snapshot has been unmounted on the slave.

Checking a File System

The **Check File System** procedure checks an unmounted file system for consistency (using the `fsck` command). See the `fsck_vxfs(1M)` manual page for more information on file system checking.

▼ To check a file system on a volume

1. Select the volume containing the file system to be checked.
2. Choose **Actions > File System > Check File System**.
3. Click **Yes** in the Check File System dialog to confirm that you want the file system checked.

▼ To check any file system

1. Select the file system to be checked.
2. Choose **Actions > Check File System**.
3. Click **Yes** in the Check File System dialog to confirm that you want the file system checked.

Notes:

- ◆ The file system must be unmounted.
- ◆ VEA assumes that `fsck` does not require user input. If user input is required, this task fails and `fsck` must be run at the command line.

Monitoring File System Capacity

The **Set Usage Threshold** procedure enables you to monitor the capacity of a file system on a mounted volume.

▼ To monitor the capacity of a file system

1. Select the file system to be monitored.
2. Choose **Actions > Set Usage Threshold**.
3. Complete the Set Usage Threshold dialog box as follows:

Monitor Capacity	Select the Monitor Capacity checkbox to initiate capacity monitoring on the file system. The File Systems Tabular View indicates if capacity monitoring is Enabled or Disabled .
Threshold Percentage	Select the Threshold Percentage value. When usage crosses this value, an alert will be logged.

Notes:

For systems other than vxfs, alert thresholds are not persistent (that is, the information is lost when the server restarts). This means that:

- ◆ File system objects on the bus reflect the threshold and enabled status for alerts.
- ◆ Once an alert is fired, the status on the bus is disabled so that the alert is not refired at every periodic rescan.

(This can be avoided by allocating a low value to Threshold Percentage.)

Viewing File System Properties

The File System Properties window contains detailed information about the selected file system. Select **File Systems** then **Properties**.

You can also view the properties associated with each mount point under the file systems.

- ◆ To view the properties of a particular file system mount point, select **File Systems**, then choose the mount point and **Properties** from the popup menu.

Using the QuickLog Feature

The QuickLog feature is available only for VxFS disk layouts 4 and 5 and requires a separate license. You must use Intent log for VxFS layout version 6.

- ◆ Creating a QuickLog Volume
- ◆ Enabling a QuickLog Device for a File System
- ◆ Disabling a QuickLog Device for a File System
- ◆ Removing a QuickLog Volume

Creating a QuickLog Volume

The **Make Log** procedure designates a volume as a QuickLog™ volume. After setting up a QuickLog volume, you can enable a QuickLog device for a file system.

▼ To designate a volume as a QuickLog volume

1. Select the volume to be used for the QuickLog device.
2. Choose **Actions > QuickLog > Make Log**
3. Complete the Make QuickLog Log dialog box as follows:

Volume Name:	Check that the correct volume name is displayed in this field.
Log Device Number:	Select the number of the log device to attach to the volume. Log devices are listed in <code>/dev/qlog</code> .
Create Cluster Quicklog checkbox (shown only if the volume is part of a cluster):	Check this box if you want the Quicklog volume to be available to the cluster.

After you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ The **Make Log** task is only available if QuickLog is installed/licensed.
- ◆ The **Make Log** task requires a log device number (1-31).
- ◆ Only VxFS file systems can be used with a QuickLog device.
- ◆ The specified volume cannot contain a file system.
- ◆ The size of the specified volume must be at least 32MB.

Enabling a QuickLog Device for a File System

The **QuickLog Enable** procedure enables a QuickLog device on a VxFS file system. Before you enable a QuickLog device, you must designate a volume as a QuickLog volume.

Note You can enable QuickLog on the file system automatically at system startup by selecting **Mount at Boot**, and specifying the appropriate options when completing the New File System dialog box (see Creating a New File System on a Volume).

▼ To enable a QuickLog device for a file system

1. Select the file system on which to enable a QuickLog device.
2. Choose **Actions > QuickLog > Enable**
3. Complete the Enable QuickLog dialog box as follows:

File System Name:	Check that the correct file system name is displayed in this field.
Log Device Number:	Select the number of the log device to be enabled for the file system. Log devices are listed in <code>/dev/qlog</code> .
Make QuickLog enable persistent (checkbox):	Check to enable QuickLog whenever the File System is mounted.

Using the QuickLog Feature

After you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ This task is only available if QuickLog is installed/licensed.
- ◆ This task requires a file system name.
- ◆ Only VxFS file systems can be used with a QuickLog device.
- ◆ The specified QuickLog device must already have a designated QuickLog volume.
- ◆ If no log device number is specified, the least-loaded device is used.

Disabling a QuickLog Device for a File System

The **QuickLog Disable** procedure disables QuickLog for a VxFS file system.

▼ To disable a QuickLog device for a file system

1. Select the file system to have QuickLog disabled.
2. Choose **Actions > QuickLog > Disable**.
3. Check that the correct file system name is displayed in the Disable QuickLog dialog box, then click **OK**.

Notes:

- ◆ The **QuickLog Disable** task is only available if QuickLog is installed/licensed.
- ◆ The **QuickLog Disable** task requires a file system name.

Removing a QuickLog Volume

The **QuickLog Remove** procedure detaches a QuickLog volume from a QuickLog device. Before you remove a QuickLog volume, you must disable QuickLog for all QuickLog-enabled file systems using that QuickLog device.

▼ To remove a QuickLog volume

1. Select the QuickLog volume to be detached.
2. Choose **Actions > QuickLog > Remove Log**.

Check that the correct volume name is displayed in the QuickLog Remove Log dialog box, then click **OK**.

Notes:

- ◆ The **QuickLog Remove** task is available only if QuickLog is installed/licensed.
- ◆ QuickLog must be disabled for all QuickLog-enabled file systems using the QuickLog device before the QuickLog volume can be detached.

Unmounting a File System from a Cluster Node

Use the **Unmount File System** procedure to unmount a VxFS file system on a selected volume in cluster (*shared*) mode.

▼ To unmount a cluster file system from a particular volume

1. Select the volume containing the cluster file system to be unmounted.
2. Choose **Actions > File System > Unmount File System...**
3. Select the nodes to be unmounted from the Unmount Cluster File System dialog box.
4. Click **Yes** in the Unmount File System dialog to confirm that you want the file system unmounted.

▼ To unmount any cluster file system

1. Select the File Systems node; then select the cluster file system to be unmounted.
2. Choose **Actions > Unmount File System...**
3. Select the nodes to be unmounted from the Unmount Cluster File System dialog box.
4. Click **Yes** in the Unmount File System dialog to confirm that you want the file system unmounted.

Removing Resource Information for a Cluster File System

The **Remove Cluster Mount** procedure removes cluster file system resource information from the cluster configuration file.

▼ **To remove a cluster file system resource information**

1. Select the file system.
2. Choose **Actions > Cluster > Remove Cluster Mount**.
3. Click **Yes** to remove the cluster file system resource information from the cluster configuration file and dissociate the shared volume from the mount point.

Notes:

- ◆ The cluster file system must be unmounted before performing the **Remove Cluster Mount** procedure.
- ◆ The **Remove Cluster Mount** procedure must be performed before changing any mount options. When remounting a previously unmounted shared volume with different mount options, the mount options will not be changed if the cluster mount point was not removed with **Remove Cluster Mount**.

Maintaining File Change Log

You can:

- Enable File Change Log
- Disable File Change Log
- Synchronize File Change Log
- Tune File Change Log

Enable File Change Log

You can enable file change log on a mounted file system.

▼ **To enable file change log**

1. Select the file system.
2. Choose **Change Log > Enable**.
3. Click **Yes** when the Confirmation message is displayed.

Disable File Change Log

You can disable file change log on a mounted file system.

▼ **To disable file change log**

1. Select the file system.
2. Choose **Change Log > Diabile**.
3. Click **Yes** when the Confirmation message is displayed.

Synchronize File Change Log

You can synchronize file change log on a mounted file system. File Change Log usually logs changes only at pre-assigned intervals. This operation forces the next structural change to the file system to be logged irrespective of the interval, and resets the time interval.

▼ **To synchronize the file change log**

1. Select the file system.
2. Choose **Change Log > Synchronize**.

Tune File Change Log

You can tune file change log on a mounted file system.

▼ **To tune the file change log**

1. Select the file system.
2. Choose **Change Log > Tune**. The Tune File Change Log dialog is displayed.
3. Complete the Tune File Change Log dialog box as follows:

Maximum size of log (bytes):	Maximum size of log. (Optional)
Keep record (seconds):	Log keep time. (Optional)
Write interval (seconds):	Logging interval. (Optional)
NetBackup is the only user for this log	Use a prolonged logging interval.

Note At least one of the parameters must be defined.

4. After you have provided all necessary information in the dialog box, click **Apply** then **OK**, or **OK**.

Maintaining Storage Checkpoints

You can:

Create a Checkpoint

Remove a Checkpoint

Mount a Checkpoint

Unmount a Checkpoint

Enable Checkpoint Quota

Disable Checkpoint Quota

Create a Checkpoint

▼ To create a checkpoint

1. Select the file system.
2. Choose **Storage Checkpoint > Create**. The Create Storage Checkpoint dialog is displayed.
3. Complete the Create Storage checkpoint dialog box as follows:

Checkpoint name:	A default checkpoint name is constructed using time zone information. You can specify your own checkpoint name, if desired.
Removable:	Default is for a removable checkpoint.
Mount:	Select this checkbox if you want the checkpoint mounted.
Mount point:	If you check Mount: , then you need to supply a mount point.
Mount details:	If you check Mount: , then mount details (Read Only , and Add to file system table) can be defined. If you are creating a checkpoint for a cluster file system, then mount details Mount and Read only can be selected for each cluster node

4. After you have provided all necessary information in the dialog box, click **Yes**.

Remove a Checkpoint

You can remove existing checkpoint; provided that it is unmounted.

▼ To remove a checkpoint

1. Select the file system.
2. Choose **Storage Checkpoint > Remove**. The remove Storage Checkpoint dialog is displayed.
3. Move the checkpoints to be removed from **All checkpoints** to **Checkpoints to remove**. (Only the checkpoints that have been unmounted are listed under **All checkpoints**.)
4. Click **OK** to remove the selected checkpoints.

Note A checkpoint must be unmounted before it can be removed. To unmount a checkpoint, select the checkpoint under the file system it is based on. Then, unmount the selected checkpoint by selecting **Actions > Unmount File Systems**.

Mount a Checkpoint

You can mount an existing checkpoint.

▼ To mount a checkpoint

1. Select the file system to provide the checkpoint.
2. Choose **Storage Checkpoint > Mount**. The Mount Storage Checkpoint dialog is displayed.
3. Complete the Mount Storage checkpoint dialog box as follows:

Checkpoint name:	A default checkpoint name is constructed using time zone information. You can specify your own checkpoint name, if desired.
Mount point:	Verify the mount point displayed is the one you need.
Mount details:	Check the required mount options (Mount and Read only for mounting checkpoint of a cluster file system; Read Only and Add to file system table for mounting checkpoint of a local file system).

Maintaining Storage Checkpoints

4. After you have provided all necessary information in the dialog box, click **Yes**.

Note Mounted Checkpoints are shown as children of the file system they are based on.

Unmount a Checkpoint

You can unmount an existing checkpoint.

▼ To unmount a checkpoint

1. Select the file system containing the checkpoint to be unmounted.
2. Choose **Storage Checkpoint > Unmount**. The Unmount Storage Checkpoint dialog is displayed.
3. Complete the Unmount Storage checkpoint dialog box as follows:

Checkpoint name:	Select the checkpoint to be unmounted.
Nodes to Unmount	Select the nodes on which the checkpoint is to be unmounted. (Move nodes from Mounted on Nodes to Nodes to Unmount .)
Remove cluster mount:	Check this checkbox to remove cluster mount from file system table.
Remove storage checkpoint	Check this checkbox to remove storage checkpoint from file system table.

4. After you have provided all necessary information in the dialog box, click **OK**.

Enable Checkpoint Quota

You can enable a checkpoint quota on a mounted file system or its hosting device (provided that the client has access to the device).

▼ To enable a checkpoint quota

1. Select the file system.
2. Choose **Storage Checkpoint > Enable Checkpoint Quotas**. The Enable Checkpoint Quotas dialog is displayed.

3. Complete the Enable Checkpoint Quotas dialog box as follows:

Hard limit:	Define the hard limit - the point at which removable checkpoints are lost. If you want, you can check the box Reuse space occupied by removable checkpoints when hard limit is reached.
Soft limit:	Define the soft limit - the point at which an alert is raised. If you want, you can check the box Allow new checkpoint creation when soft limit is reached.

Note The hard limit for the checkpoint cannot be less than the soft limit.

4. After you have provided all necessary information in the dialog box, click **Yes**.

Disable Checkpoint Quota

You can disable a checkpoint quota on a mounted file system or its hosting device (provided that the client has access to the device).

▼ To disable a checkpoint quota

1. Select the file system.
2. Choose **Storage Checkpoint > Disable Checkpoint Quotas**. The Disable Checkpoint Quotas confirmation dialog is displayed.
3. Click **Yes** to confirm.

Using Multi-Component File Systems

VxFS supports file systems spanning multiple devices (Multi-Component Support, or MDS). Such systems are called MDS file systems, and can be regarded as residing on a VxVM device - a VxVM volume set.

- ◆ Setting up an MDS System
- ◆ Adding Volumes
- ◆ Removing Volumes
- ◆ Encapsulating Volumes

Setting up an MDS System

Typically setting up an MDS file system involves:

- ◆ Creating a Volume Set (see Chapter 5)
- ◆ Creating the MDS File System (see [Creating a New File System on a Volume Set](#))
- ◆ Mounting the Volume Set

Note If an MDF is selected, the sub file system objects are show in a tab view in the table.

Mounting the Volume Set

You can mount the file system on the volume set as follows:

▼ **To mount a file system**

1. Select the volume set.
2. Choose **Actions > File System > Mount File System**. The Mount File System dialog appears (see [Mounting a File System](#)).

Adding Volumes

Note Before adding a volume to an MDS file system, the volume must already have been added into volume set.

▼ To add volumes to an MDS file system

1. Select the MDS file system.
2. Choose **Actions > Multi-Component File System > Add Volume**. The **Add Volume to File System** dialog appears.
3. Move the required volumes from the **Volumes not in File System** area to the **Selected volumes** area.
4. Click **OK**.

Note When adding a volume to a volume set, the volume is also added to an existing MDS file system.

Removing Volumes

Note You must remove sub file systems manually. This operation can take a long time to complete (it is asynchronous). The data needs to be moved off the subvolume to another volume in the volume set.

▼ To remove a volume from an MDS file system

1. Select the MDS file system.
2. Choose **Actions > Multi-Component File System > Remove Volume**. The **Remove Volume** dialog appears.
3. Select the volume to be removed.
4. Click **OK**.

Encapsulating Volumes

▼ To encapsulate a volume in an MDS file system

1. Select the MDS file system.
2. Choose **Actions > Multi-Component File System > Encapsulate Volume**. The **Encapsulate Volume** dialog appears.

Set Intent Log Options

3. Select the volume to be encapsulated.
4. You can define a new file name, if you want (default file name is same as volume name).
5. Click **Yes**.

Set Intent Log Options

Intent log is available on VxFS layout version 6. (Use Quicklog for layout versions 4 and 5.)

Note VxFS layout version 6 is the default when a file system is created using the GUI.

For MDS file systems, this is a replacement for quicklog.

▼ To set intent log options

1. Select the file system.
2. Choose **Set Intent Log Options**. The Intent Log Options dialog is displayed.
3. Complete the Intent Log Options dialog box as follows:

Intent log size:	Size of the intent log. It must be greater than 256K bytes, and less than 262144K bytes.
Device to use:	Name of volume set to use for intent logging.

4. After you have provided all necessary information in the dialog box, click **Yes**.

Using the Quality of Storage Service Feature 8

The VERITAS File System Quality of Storage Service (QoSS) feature allows you to map more than one device into a single file system. (Refer to the *VERITAS File System Administrator's Guide* for background information on how to configure file relocation policies, and how to use QoSS command-line utilities.)

Note VERITAS Quality of Storage Service is a licensed feature. You must purchase a separate license key to use QoSS, and the `VRTSfppm` package must be installed for the VERITAS Quality of Storage Service feature to operate. (See the *VERITAS File System Release Notes* for current product information.)

QoSS Task Roadmap

The following table provides a summary of common QoSS tasks.

To perform this task ...	Use this procedure ...
Configure Relocation Policies	Configuring Relocation Policies
Analyze a file system	Analyzing a File System
Relocate files	Relocating Files
Select Files	Selecting Files Based on Maximum Age

Configuring Relocation Policies

If a file system has multiple component devices, you must configure relocation policies to define the files to move and the physical devices on which to locate the files. Policies for a file are based on the following criteria:

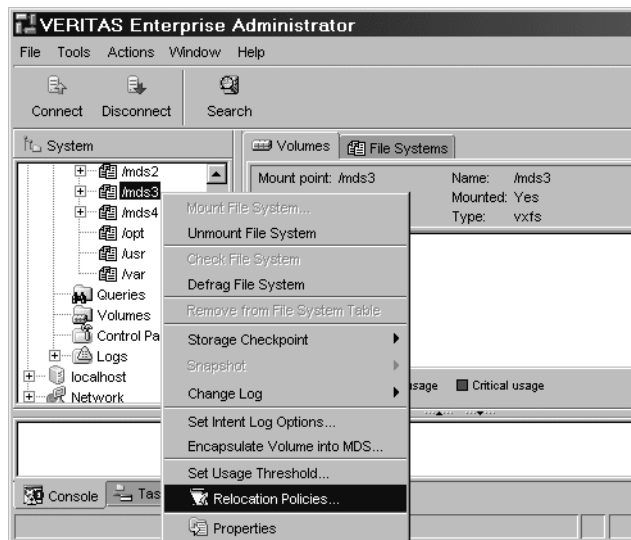
- ◆ Age, which is the time since it was last accessed or last modified
- ◆ Size
- ◆ Location in a directory structure
- ◆ Name pattern, such as the similarly named output files of an application or a common extension such as *.gif
- ◆ Source component volume
- ◆ Destination component volume

Viewing Relocation Policies and Modifying Policy Rules

▼ To view relocation policies or modify policy rules

1. Right-click one or more file systems that have multiple component devices.
2. Select **Relocation Policies** from the popup menu.

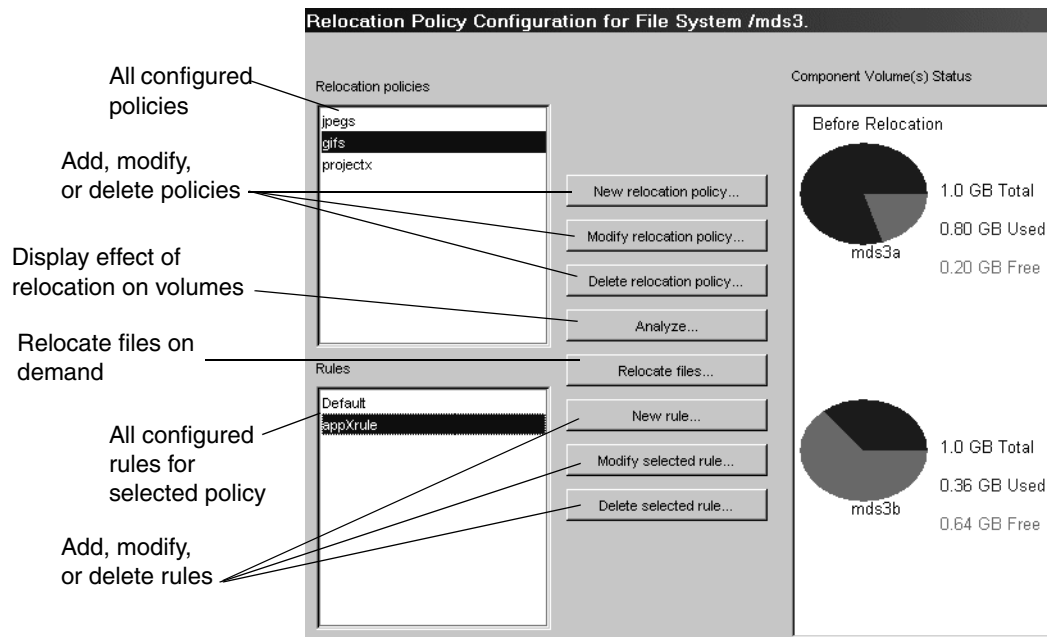
Selecting Relocation Policies



The Relocation Policies Configuration dialog appears.

The top left pane displays all the policies that have been configured for the file system. The bottom left pane displays all the rules for the relocation policy selected in the top list. The panes on the right display the disk space available and used for the component volumes in the file system.

Relocation Policies Configuration Dialog



3. From the Relocation Policies Configuration screen, you can do any of the following:
 - a. Click **New relocation policy** to launch the Relocation Policy Wizard and use it to create a new policy.
 - b. Click **Modify relocation policy** to launch the Relocation Policy Wizard and use it to modify an existing policy. You can add rules to a policy, delete source or destination component volumes from a policy, and so on.
 - c. Click **Delete relocation policy** to remove the selected policy. You must confirm your choice before the policy will be deleted.

Configuring Relocation Policies

- d. Click **Analyze** to display the effect of relocation policies on a volume. The results are displayed in the pie chart on the far right. Clicking this button does not relocate files; it only displays how disk space would change if files were relocated according to a specific policy. For detailed information, see “To experiment with relocation policies using the Analyze button” on page 204.
- e. Click **Relocate files** to relocate all eligible files immediately. Usually files are relocated according to a predetermined schedule. For more information, see “To relocate files immediately” on page 207.
- f. Click **New rule** to create a new rule that will be associated with an existing policy that is selected from the list in the top pane.
- g. Click **Modify selected rule** to modify a rule associated with an existing policy that is selected in list in the top pane.
- h. Click **Delete selected rule** to delete a rule that is associated with an existing policy that is selected in list in the top pane

Note The **Default** rule can not be deleted because it is always associated with the policy.

- i. Click the **Apply** button to save your changes and leave the dialog open for further modifications.
- j. Click the **Close** button and confirm your choice to save your changes and close the dialog.

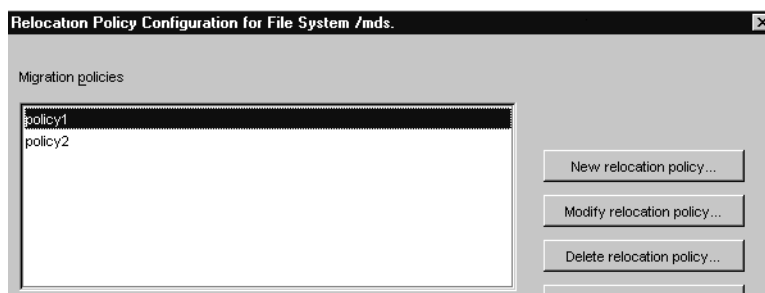
Creating and Modifying Relocation Policies

▼ To create or modify a relocation policy

1. Right-click one or more file systems that have multiple component devices.
2. Select **Relocation Policies** from the popup menu.
3. From the Relocation Policies Configuration dialog, click the **New relocation policy** button to launch the Relocation Policy Wizard.

To modify an existing policy, select it and click the **Modify relocation policy** button. You cannot modify the name of a policy using the wizard.

Relocation Policy Configuration Wizard



The Relocation Policy Wizard launches.

4. After reading the Welcome screen, click **Next**.
The Source Component Volume screen appears.
5. On the Source Component Volume screen, specify the policy name and select the source component volume or volumes the policy will manage.

Configuring Relocation Policies

Source Component Volume Screen

Component Volume	Usage	Capacity	Free Space
mds3a	80 %	1.00 GB	0.20 GB
mds3b	36 %	1.00 GB	0.64 GB
mds3c	19 %	1.00 GB	0.80 GB

The source component volume has files that you want to manage by relocating them when they match relocation policy rules. By selecting a volume, you configure the file system sweeping utility to search it for files that can be relocated based on your configuration.

Click **Next**.

- On the Destination Component Volume screen, select the component volume or volumes to which the data will be moved from the source component volume or volumes.

Destination Component Volume Screen

Component Volume	Usage	Capacity	Free Space
mds3b	36 %	1.00 GB	0.64 GB
mds3c	19 %	1.00 GB	0.80 GB

The destination component volume is the location for files selected for relocation. By choosing a volume, you configure the file system sweeping utility to relocate files it selects to that component volume.

Click **Next**.

- On the Relocation Rules screen you can add, modify, and remove rules for a relocation policy.

Relocation Rules Screen

grady - New Relocation Policy

Relocation Rules
Specify the relocation rules. Files that match these rules are eligible to be relocated.

Rule Name	Source Vol...	Destination	Action	File Type	Path	Larger (KB)	Smaller (KB)	Accessed (days)
Default	mds3a	mds3b	Include	*	/mds3	10000	1000000	30
appXrule1	mds3a	mds3b	Include	*	/mds3	10000	1000000	20
logs	mds3a	mds3c	Include	*.log	/mds3	10000	1000000	14

Column headings

Arrow buttons

New Rule... Modify Selected Rule... Delete Selected Rule... ▲ Move rule up... ▼ Move rule down...

You can use the column headings to sort the information on this screen. Click a column heading to organize the information by that criteria. For example, to display all of the files from the same source component volume together in the list, click the **Source Volume** heading.

Use the arrow buttons to move a selected row up or down.

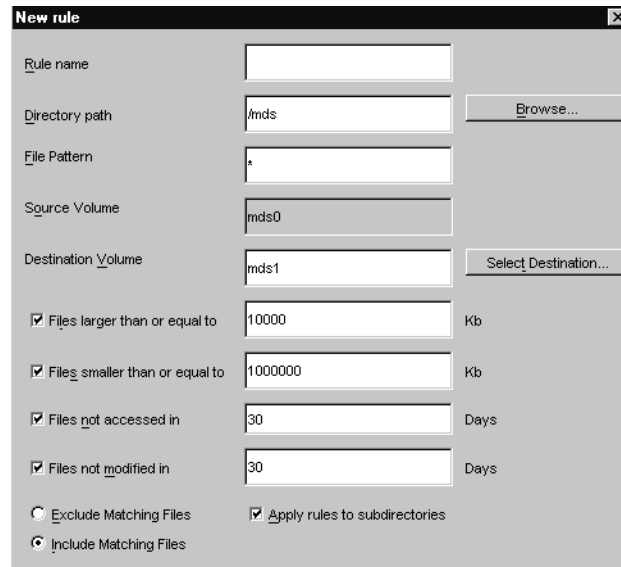
Click the **Delete Selected Rule** button to delete the selected rule (you must confirm your choice).

Note The **Default** rule can not be deleted because it is always associated with the policy. You can change the name Default through the language-dependent Java properties file for the GUI extension.

Click the **New Rule** button to create a new rule. The Relocation Rules dialog appears.

Configuring Relocation Policies

Relocation Rules Dialog



New rule

Rule name:

Directory path:

File Pattern:

Source Volume:

Destination Volume:

Files larger than or equal to Kb

Files smaller than or equal to Kb

Files not accessed in Days

Files not modified in Days

Exclude Matching Files Apply rules to subdirectories

Include Matching Files


- a. The only required criteria for a rule is a name that is unique among the other rules for relocation policies. The relocation rule you name will use the criteria you select to determine which files to select for relocation.

Note A file must meet all of the criteria you specify in a rule in order to be eligible for relocation. For example, if you select both a directory path and file pattern, the sweep of the file system will consider only files that meet both conditions.

- b. Specify a directory path for files to which the rule will apply.

To select a directory from the source component volume or volumes from a list of possible directories, click **Browse**. The Select Directory dialog appears.

Select Directory Dialog



Select a directory or directories for this rule and click **OK**.

- c. Specify the destination component volume to which this rule will apply. This is the component volume to which files will be relocated.

To select a destination component volume or volumes from a list of possible volumes, click **Select Destination**. The Select Volume dialog appears.

Select Volume Dialog


Component Volume	Usage	Capacity	Free Space
mds3b	36 %	1.00 GB	0.64 GB
mds3c	19 %	1.00 GB	0.80 GB

Select a component volume or volumes for this rule and click **OK**.

- d. After returning to the Relocation Rules dialog, you can specify how the rule will be used:

- ◆ Exclude files when the rule is applied
- ◆ Include files when the rule is applied

Specifying how a Rule is to be Applied



- e. To apply this rule to all subdirectories under those the rule causes to be selected, select **Apply rule to subdirectories**.

Configuring Relocation Policies

- f. To add, modify or remove file selection criteria from the rule, toggle the checkmark next to the criterion either on or off, as appropriate. If you remove a checkmark, the criterion is not considered in selecting files to relocate.

Adding, Modifying or Removing File Selection Criteria

<input checked="" type="checkbox"/> Files larger than or equal to	<input type="text" value="10000"/>	Kb
<input checked="" type="checkbox"/> Files smaller than or equal to	<input type="text" value="1000000"/>	Kb
<input checked="" type="checkbox"/> Files not accessed in	<input type="text" value="30"/>	Days
<input checked="" type="checkbox"/> Files not modified in	<input type="text" value="30"/>	Days

You can set criteria in the following areas:

File size

- ◆ The **Files larger than or equal to** field specifies that files of the given size and greater can be relocated. To set the criteria, click the checkbox and enter a value to the right.
- ◆ The **Files smaller than or equal to** field specifies that files of the given size and smaller can be relocated. To set the criteria, click the checkbox and enter a value on the right.

File age

- ◆ The **Files not accessed in** field specifies that files accessed in the given period of time are not relocated.
- ◆ The **Files not modified in** field specifies that files modified in the given period of time are not relocated.

Click **Next**.

The Policy Configuration Summary screen appears.

Policy Configuration Summary Screen

Policy Configuration Summary

You have selected the following relocation policy configuration:

Migration policy name: policy3

Rule Name: Default
Source component volume: mds0
Destination component volume: mds1
Action: Include
Apply rules to subdirectories: Yes
Directory path: /mds
File Pattern: *
Files larger than: 10000 Kb
Files smaller than: 1000000 Kb
Files not accessed in: 30 Days
Files not modified in: 30 Days

Rule Name: jpegs
Source component volume: mds0

< Back Finish Cancel Help

8. If the screen summarizes the policy as you want to configure it, click **Finish** to configure the policy.

If you wish to make changes, click **Back**, modify the appropriate rules or policies, and click **Next** until you return to the Policy Configuration Summary screen. Click **Finish** to configure the policy.


Analyzing a File System

One method for learning how relocation policies affect a file system is to experiment with scenarios based on different policies, such as how configurations affect the disk space on volumes in a file system.

▼ To experiment with relocation policies using the Analyze button

1. In VEA, right-click one or more file systems that have multiple component devices.
2. Select **Relocation Policies** from the popup menu.

Selecting Relocation Policies



The Relocation Policies Configuration dialog appears.

3. Determine the relocation policy that you want to test and select it from the Relocation policies pane. You can also test how relocating files according to all policies will affect the disk space on component volumes.
4. To display the effect of relocation policies on a volume, click **Analyze**. The results are displayed in the pie chart on the far right.

Relocation Policy Analysis Screen

Relocation Policy Configuration for File System /mds3.

Relocation policies

- jpegs
- gifs**
- projectx

Rules

- Default
- appXrule**

Component Volume(s) Status

Volume	Total	Used	Free	Relocated Data	New Data
mds3b	1.0 GB	0.36 GB	0.64 GB	0.0 GB	0.0 GB
mds3c	1.0 GB	0.20 GB	0.80 GB	0.0 GB	0.0 GB

5. A confirmation dialog prompts you to confirm whether you want to analyze the file system based on the policy that is selected in the Relocation policies pane or based on all policies. Click the appropriate radio button and **OK**.
6. After you have confirmed your choice of policies, an Analysis Status window appears that shows the total number of files that will be analyzed, how many have been processed so far, how long the analysis is expected to take, and how much time it has taken.

Analysis Status Window

Analysis Status

76%

Total Files: 9058
 Processed Files: 6900
 Estimated Time (HH:MM:SS): 00:00:39
 Elapsed Time (HH:MM:SS): 00:02:07

Cancel

Analyzing a File System

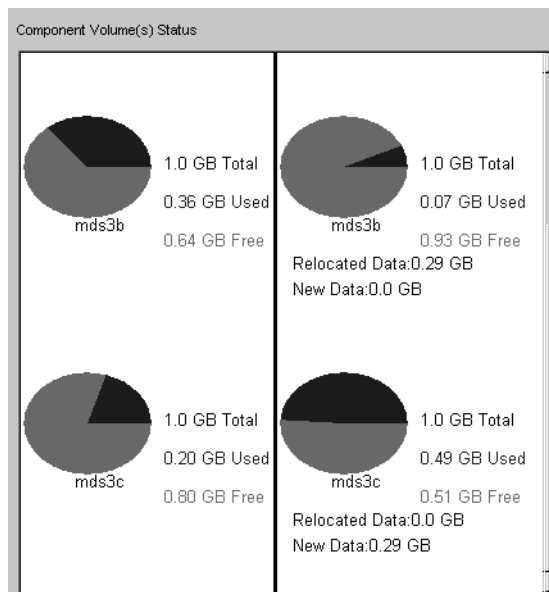
7. After the status bar at the top of the Analysis Status window has reached 100%, click **Close** and examine the Component Volume(s) Status to see where files will move within the file system.

Clicking the **Analyze** button does not relocate files; it only displays how disk space would change if files were relocated according to a specific policy.

After you have analyzed a policy, you can relocate files based on the results by click **Relocate files** and confirming your choice when prompted.

You can select different policies and see corresponding results.

Relocation Policy Results Screen



Relocating Files

After you have configured policies, you can use them to select and relocate files.

▼ To relocate files immediately

1. Right-click one or more file systems that have multiple component devices.
2. Select **Start Relocations** from the popup menu.

Selecting Files Based on Maximum Age


The Configuration Policy Wizard lets you relocate files based on minimum file size and maximum file size, but only on minimum file age.

It is useful to select files based on maximum age, that is, files that were recently accessed or modified. For example, you can add a new component volume to an existing volume set, and the newest component volume is on more modern hardware with improved performance. You may want to relocate the most recently accessed files to this new component volume.

To select files on the basis of a maximum age, use two rules within one policy. The first policy rule can have any file size or name criteria, but it has no age criteria at all. The second rule has no file size or name criteria, but it does have a minimum age and is marked in the **Relocation Rules** dialog to **Exclude matching files**.

When QoSS evaluates files, it considers all policies and rules. If a file meets the criteria of a rule, it is relocated until all other rules have been evaluated. If one rule includes a file, but another rule excludes the same file, the file is not relocated.

So in this example, one rule selects files of any age, possibly constrained by criteria other than age. The second rule selects files older than a minimum age and excludes them from relocation. Because the first rule is inclusive and the second rule is exclusive, the effect is to stop relocation of older files and to relocate recent files.



Selecting Files Based on Maximum Age



Device-Related Tasks

9

Device-related tasks are handled by the Device Discovery layer (DDL). It supports device discovery and allows dynamic addition of support for disk arrays for multipathing.

Procedures are provided for:

- ◆ Maintaining the List of Foreign Devices
- ◆ Maintaining Disks
- ◆ Enabling and Disabling Paths
- ◆ Displaying Lists of Supported Arrays

The Device Task Roadmap shows a summary of common DDL tasks.

Device Task Roadmap

The following table provides a summary of common device tasks.

To perform this task ...	Use this procedure...
Maintaining Foreign Device List	
Add a foreign device entry to the <code>/etc/vx/darecs</code> file.	Add Foreign Devices
Remove a foreign device entry, or directory, from the <code>/etc/vx/darecs</code> file.	Remove Foreign Devices
List information about a foreign device.	Display Information for a Foreign Device
Maintaining Paths	
Enable/Disable a path	Enabling and Disabling Paths
Maintaining Disks	
Display a list of supported arrays for a host.	Displaying Lists of Supported Arrays

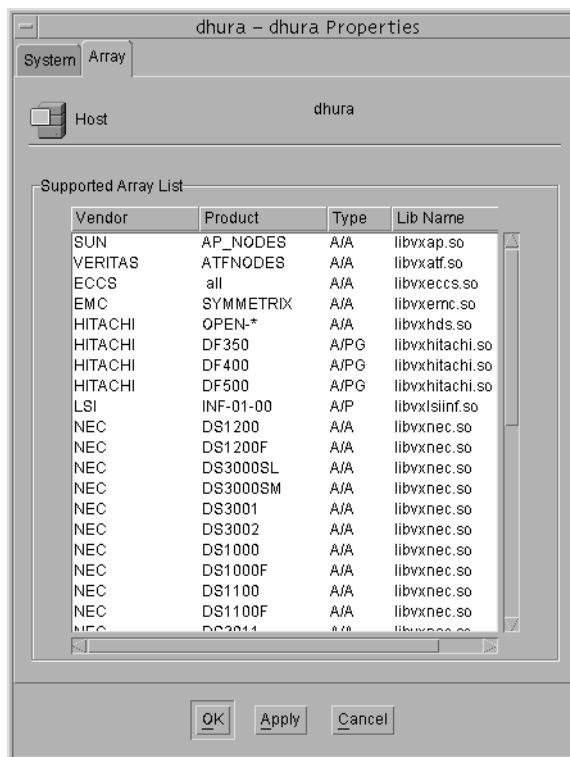
Displaying Lists of Supported Arrays

Use the following procedure to display a list of supported arrays for a host.

▼ **To display a list of supported arrays**

1. Right-click on the host for which the supported arrays are to be displayed.
2. Choose **Properties** from the displayed sub-menu.
3. Click on the **Arrays** tab
4. The arrays supported by this host is displayed.
5. **OK** to exit.

List of Supported Arrays



Enabling and Disabling Paths

Use the following procedure to enable or disable a path to a controller.

▼ **To enable or disable a path**

1. Select the controller node in the left pane. The controller table view appears.
2. Select the controller containing the paths to be enabled or disabled.
3. Select the Paths tab from the table view in the right pane. The path name and status (enabled/disabled) is displayed.
4. Select the paths to be enabled or disabled.
5. To enable a path, select **Actions > Enable Path** , to disable a path, select **Actions > Disable Path**. A confirmation message is displayed.
6. Choose **OK** to confirm that you want to enable or disable the paths shown.

Maintaining Disks

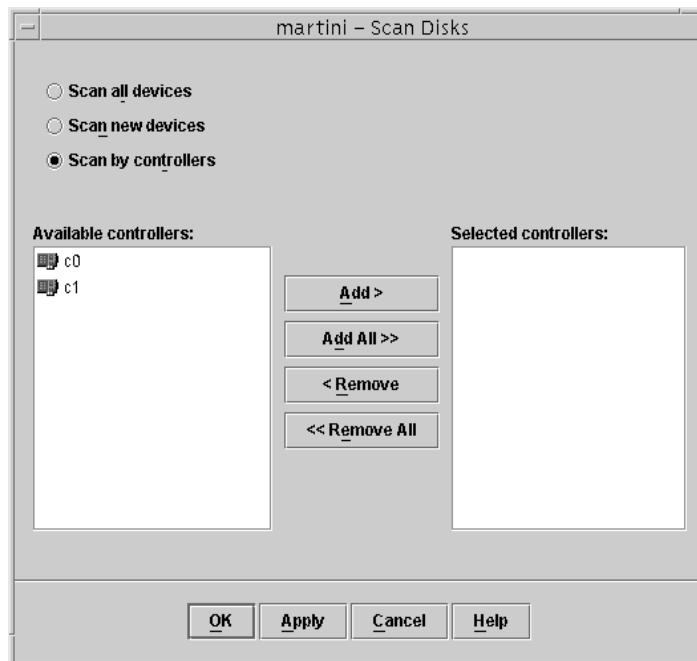
You can initiate a rescanning of the devices in the device tree. Use the following procedure to scan devices in the operating system device tree and to trigger a Dynamic MultiPathing (DMP) reconfiguration.

Note This operation will take a long time to complete.

▼ To scan disks

1. Select the disks node in the left pane.
2. Select Actions > Scan Disks, a Scan Disk dialog box is displayed. You can now choose to:
 - ◆ Scan all devices
 - ◆ Perform a partial discovery of devices (Scan new devices)
 - ◆ Scan by controllers

Scan Disks Dialog



Maintaining Disks

3. If you choose to scan by controllers, you need to select at least one controller to scan (move from **Available controllers:** to **Selected controllers:**).
4. Select **OK**; then choose **Yes** to confirm that you want to continue with the disk scan.

Maintaining the List of Foreign Devices

Procedures are provided for you to maintain the list of foreign devices - those devices that are not auto configurable by VxVM. You can:

- ◆ Display Information for a Foreign Device
- ◆ Add Foreign Devices
- ◆ Remove Foreign Devices

Display Information for a Foreign Device

Use the following procedure to list information about the device and option information for the list of foreign devices in the `/etc/vx/darecs` file.

▼ To display information for a foreign device

1. Select the host for which the devices are to be displayed.
2. Choose **Actions > Foreign Device > List**. The List Foreign Device information appears.
3. **OK** to exit.

Add Foreign Devices

Use the following procedures to add a specific device, or a set of devices, to the list of devices in the `/etc/vx/darecs` file.

▼ To add a specific device

1. Select the host where the devices are to be added.
2. Choose **Actions > Foreign Device > Add**. The Add Foreign Devices dialog appears.

Maintaining the List of Foreign Devices

Add Foreign Devices Dialog



3. Type the absolute device path name in the **Device Directory/Path** field.
4. Select the type (**Block and char path** or **Path**), as appropriate.

Block and char path	If you use block path, then you also have to specify character path; for example, /dev/dsk/c2t0d0s2 /dev/rdisk/c2t0d0s2
Path	Example: /dev/dsk/c2t0d0s2

5. Select the options for suppressing duplicate disks

Suppress autoconfig disks	Suppress the auto discovered entries in case of duplicates (that is, keep the foreign disk entries)
Suppress foreign disks	Suppress the foreign entries in case of duplicates (that is, keep the auto discovered entries)
Do not suppress any disk	Duplicate entries are not suppressed.

6. Select **OK**. The device is now added to the `/etc/vx/darecs` file.

▼ **To add a set of devices**

1. Choose **Actions > Foreign Device > Add**. The Add Foreign Devices dialog appears.
2. Type the name of the directory containing the set of devices to be added in the **Device Directory/Path** field.
3. Select the type (**Block and char directory** or **Directory**), as appropriate.

Block and char directory	If you specify a block directory, then you also have to specify a character directory; for example, <code>/dev/dsk /dev/rdisk</code>
Directory	Example: <code>/dev/dsk</code>

4. Select the options for suppressing duplicate disks

Suppress autoconfig disks	Suppress the auto discovered entries in case of duplicates (that is, keep the foreign disk entries)
Suppress foreign disks	Suppress the foreign entries in case of duplicates (that is, keep the auto discovered entries)
Do not suppress any disk	Duplicate entries are not suppressed.

5. Select **Apply** then **Cancel**. The device is now added to the `/etc/vx/darecs` file.

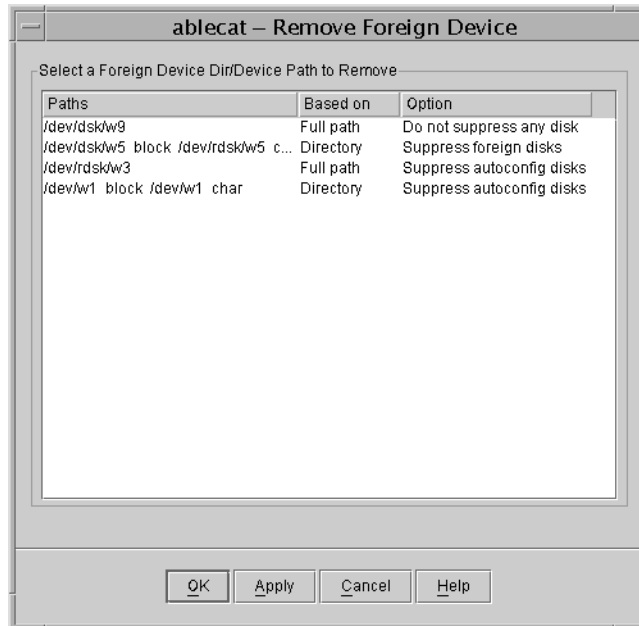
Remove Foreign Devices


Use the following procedure to remove a specific device, or a directory from the list of devices in the `/etc/vx/darecs` file.

▼ To remove foreign devices



1. Select the host from where the devices are to be removed.
2. Choose **Actions > Foreign Device > Remove**. The Remove Foreign Device dialog appears.
3. Select the directory or device path to be removed.
4. Click **OK** to apply the changes. The confirmation message is then displayed.
5. Choose **Yes** to confirm removal.
6. Click **OK** to exit if you are making no more changes. If you clicked **Apply** in step 4, then click **Cancel** to exit if you are making no more changes.

Remove Foreign Devices Dialog





Maintaining the List of Foreign Devices



Using Intelligent Storage Provisioning

10

The VERITAS Intelligent Storage Provisioning (*ISP*) option allows you to create virtual storage devices (*application volumes*) for use in organizing and managing physical storage.

ISP uses the requirements (*capabilities*) and rules that you specify for the application volumes to create volumes from the available storage. You then use these application volumes with VEA in the same way as you would traditional volumes.

ISP also interacts with other VERITAS components, where these are available, such as the VERITAS Array Integration Layer (VAIL) and the SAN Access Layer (SAL), to take appropriate actions; as for example, when configuring intelligent storage in a SAN environment.

Benefits of Using ISP

The following list summarizes the main benefits that *ISP* provides over the existing storage allocation features:

- ◆ Storage is automatically allocated based on abstract requirements such as the desired capabilities of a volume.
- ◆ Prefabricated capabilities that are provided by vendor-specific features of intelligent storage arrays can be encoded as storage attributes, and used to allocate storage.
- ◆ All *ISP* operations preserve the original intent of the volumes. There is no possibility that operations such as grow, evacuate, add mirror, or add column can accidentally degrade the reliability or performance capabilities of a volume.
- ◆ *ISP* is SAN-aware and understands SAN attributes. It is also capable of using VAIL to learn the capabilities of LUNs.

ISP provides you with GUI interfaces for:

- ◆ Managing User Templates
- ◆ Managing Storage Pools
- ◆ Annotating a Disk

The *ISP* Task Roadmap shows a summary of common *ISP* tasks.

ISP Task Roadmap

The following table provides a summary of common ISP tasks.

To perform this task ...	Use this procedure...
Manage User Templates	
Create user templates	Create a User Template
Delete user templates	Delete a User Template
Modifying user templates	Edit a User Template
Manage Storage Pools	
Organize storage by grouping into storage pools	Creating a Storage Pool Creating a Storage Pool
Create storage pool	Creating a Storage Pool
Add or remove storage pool disks	Adding and Removing Storage Pool Disks
Annotate Disks	
Adding user attributes	Add a New User Attribute
Editing user attributes	Modify a User Attribute
Removing user attributes	Remove a User Attribute
Sorting user attributes	Sort Disks by Attribute Name

Managing User Templates

A volume template (or *template*) is a meaningful collection of rules that provide a capability. A template can specify one or more capabilities that a volume created using the template may have, and consists of a collection of storage selection and layout rules. For example, a template may allow you to create a volume that is able to tolerate the failure of a certain number of controllers, or that has a certain number of copies of the volume data.

When creating a volume, it is easier to specify its desired capabilities than to specify the precise layout of the volume on the available storage. ISP selects the appropriate templates and uses them to create a volume with the desired capabilities.

If you specify parameter values for a volume's capabilities, the rules that are defined within the chosen template use these values when selecting and laying out storage. If not specified, the default parameter values for a volume's capabilities are assumed by the template.

When using ISP to create an application volume with a specified set of capabilities, you can usually enter values for the parameters that modify these capabilities. For example, you can select the degree of redundancy, fault tolerance, or performance that you require from the volume. For convenience in creating volumes, you can set up user templates to store commonly used sets of capabilities and parameter values. You can then specify the name of a user template instead of the name of a capability when creating an application volume.

Application Volumes

An application volume is created by ISP, and then exported for use by an application such as a database or file system.

Intent

The intent of a volume is a conceptualization of its purpose as defined by its characteristics and implemented by a template. ISP attempts to preserve the intent of a volume whenever the volume is reconfigured, resized, or relocated. Intent preservation automatically conserves capabilities such as reliability and performance, and observes additional rules such as allocating storage based on confinement and exclusion specifications.

Template Set

A template set consists of related capabilities and templates that have been collected together for convenience. Associating a template set with a storage pool is equivalent to associating each of its member templates separately with the storage pool. See Template Sets, in Appendix B, for details of the template sets that are provided.

User Template

A user template (or user-defined template) defines an arbitrary collection of capabilities, templates and rules to which you want volumes of a certain type to conform. For example, you might want all volumes that you create to store database tables to share the same reliability and performance capabilities, and also that they only be allocated from a restricted set of storage.

It is useful to create user templates if you regularly create volumes with similar capabilities.

A user template consists of a list of capabilities and rules. Associated with each capability is a list of parameter-value pairs that define default parameter values for the capability. If a parameter is defined as a constant, you cannot override this value. Otherwise, you can change the value of the parameter when creating a volume.

You can recreate the same type of ISP volume (a volume having the same capabilities, parameters and rules) by storing its configuration as a *user template*. The user template can then be used in subsequent volume creations.

Use the User Template Management feature to:

- ◆ Create a User Template
- ◆ Delete a User Template
- ◆ Edit a User Template

Create a User Template

To create a user template, you:

- ◆ Create the User Template
- ◆ Select Capabilities
- ◆ Set Configuration Rules
- ◆ Check your selections on the Summary Page

Create the User Template

Use the following procedure to create a user template:

▼ **To create a user template**

1. Select **User Template** from the left pane of the VEA interface.
2. Select **Actions > New User Template**. There is a short interval while status information is downloaded. The User Template Creation Wizard then appears.
3. Click **Next>** to continue. The User Template Creation Wizard dialog box appears.
4. Complete the User Template Creation Wizard dialog box as follows:

User Template Name:	Enter the name for the user template. The name must begin with either underscore or an alphabetic character.
User Template Description:	Enter descriptive information (optional).

5. Click **Next>** to continue. The Select Capabilities dialog box appears.

Select Capabilities

A capability is a feature that is provided by a volume. For example, a volume may exhibit capabilities such as performance and reliability to various degrees. Each type of capability is defined as a set of rules.

You can select the capabilities for your template and define the parameter values for these capabilities.

Note To create application volumes successfully, the appropriate licenses must be present on your system. For example, you need a full VERITAS Volume Manager™ license and a VERITAS FlashSnap™ license to use the instant snapshot feature. Vendors of disk arrays may also provide capabilities that require special licenses for certain features of their hardware.

▼ To select a capability

1. Click on the required top-level capability (the *root capability*) in the left pane. The tree displays a *capability tree* derived from the root capability. A description of the selected root capability is also displayed, beneath the capability tree.
2. Select a capability.
3. Define the ISP parameters (select name and define a value) for the selected capability in the right pane.
4. Repeat step 2 and step 3 until you have selected and defined parameters for all the capabilities you need.
5. You can restore the original values by clicking **Restore Defaults**.

Note If the Show rules page preference (refer to Show Rules Page) is set, the Set Configuration Rules dialog box appears.

Set Configuration Rules

A rule is a statement that specifies how a volume is to be created. A rule can define selection of storage. Rules are usually gathered together as templates for creating volumes, rather than being specified individually.

Storage Selection Rules

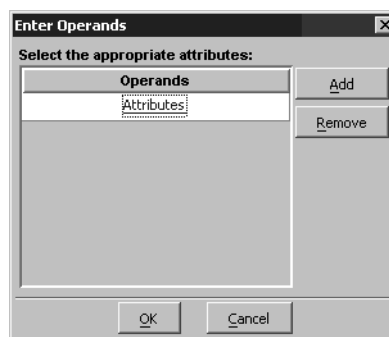
Storage selection rules specify what storage can be used to create volumes. The Preferences Show rules page box (refer to Setting Preferences in Chapter 2) must be checked if you want to enable the rules page. By default, the rules page is disabled, and ISP applies the default rules for volume and template creation, and snapshot and transform wizards.

Use the Rule Configuration Dialog box to add or remove rules defining the devices to be selected.

▼ To configure rules

1. Click on the rule cell in the left column. The Storage Selection Rules are then shown.
2. Select the appropriate rule from the list.
3. Click on the attribute cell in the right column. A button with ellipsis appears on the right side of this cell.
4. Click on the ellipsis button. The Enter Operands window is displayed.

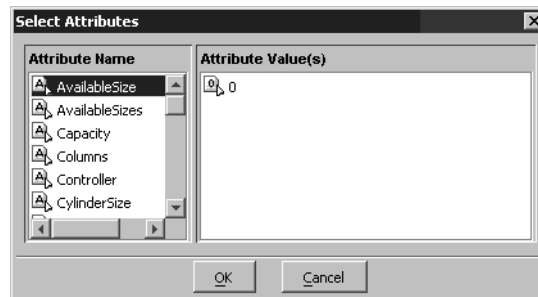
Enter Operands Window



5. Click on **Attributes**. The Select Attributes dialog appears.

Create a User Template

Select Attributes Dialog



6. Select the attribute (**Attribute Name**) and value (**Attribute Value**) for this rule.
7. Click on **OK**. You are returned to the Enter Operands dialog.
8. Click on the **Add** button to select another attribute for the rule. (You can delete a selected attribute by clicking on **Remove**.)
9. Click **OK** to return to the rules page.
10. To specify a new rule, click on the **Add** button in the Rule Configuration Dialog box.
11. When you have finished defining the rules, press **Next>** to bring up the **Summary** Page.

(To delete a rule, click **Remove** in the Rules Configuration Dialog box.)

Summary Page

The **Summary page** displays information on the name and description of the user template, and its selected capabilities and rules.

Click **Finish**. If the user template creation is successful, the message `User template created successfully`, is displayed and the wizard exits. If an error occurs while creating the user template, a message along with the error string from the provider is displayed and the wizard exits.

Performance

Performance specifies the capabilities of a volume based on factors such as number of columns in striped volumes, stripe unit size, preferred characteristics of storage such as using prefabricated disks that are configured in hardware.

Reliability

Reliability specifies the level of redundancy that is required from a volume as a capability. Very reliable volumes have a high degree of redundancy. For example, a very reliable volume could be configured as a software mirrored volume built upon underlying prefabricated RAID-5 capable LUNs that are configured in the enclosure hardware. A medium-level redundant volume could be a simple 2-way mirror or RAID-5 volume configured either in software or on a suitable LUN within a single enclosure.

Edit a User Template

Use the following procedure to change the settings on an existing user template:

▼ **To modify a user template**

1. Select **User Template** from the left pane of the VEA interface.
2. Right click the user template you want to edit from the right pane.
3. Select **Edit**. There is a short interval while status information is downloaded. The Edit User Template Wizard then appears.
4. Click **Next>** to continue. The User Template Information dialog is displayed. Enter a description of the new template under **User Template Description**.
5. Click **Next>** to continue. The **Select Capabilities** dialog is displayed. Even if a user template doesn't have any capabilities defined the capability page appears. Only if there are no capabilities defined in the configuration database will the capability page not appear.
6. Modify the selection of capabilities, if necessary. Click on the required top-level capability (the *root capability*) in the left pane. A description of the selected root capability is displayed beneath the capability tree.
7. Define the ISP parameters for the selected capability in the right pane.
8. Repeat step 6 and step 7 until you have selected and defined parameters for all the capabilities you need.
9. Click **Next>** to continue. The **Select Rules** dialog is displayed. See Storage Selection Rules for more information.
10. Click **Finish** to continue.
11. The Summary page appears.
12. Check your changes.
13. Click **Finish**. If the user template was edited successfully, a message `User template edited successfully`, is displayed and wizard exits. If an error occurs while editing the user template, a message along with the error string from the provider is displayed and the wizard exits.

Delete a User Template

Use the following procedure to delete a user template:

▼ **To delete a user template**

1. Select **User Template** from the left pane of the VEA interface.
2. Right click the user template you want to delete from the right pane. A dialog box appears asking if you are sure you are sure you want to delete the selected user template.
3. Click **Yes** to delete the user template.

Managing Storage Pools

Storage pools provide a means for organizing storage in groups, so that storage with similar characteristics is grouped for creating application volumes. A storage pool is defined within a disk group in VEA for use by ISP, and is characterized by the volume templates and LUNs it contains, and policies defined on them.

Note The disks that you assign to a storage pool must have already been initialized for use, and must belong to the disk group in which you are creating the storage pool.

The first storage pool that you create in a disk group is a Data Pool that contains application volumes. Any storage pools that you subsequently create in the disk group are Clone Pools that can be used to hold full-sized instant snapshots of the volumes in the data storage pool. You need only place such snapshots in a separate clone storage pool if they need to be created using different templates from their parent volumes, or if they are to be moved into a different disk group.

For convenience, storage pool definitions are provided that include a number of associated templates that can be used for different purposes. See “Storage Pools” on page 284 for details of the storage pool definitions that are provided.

Note Only full-sized and space-optimized instant snapshots are supported for use with ISP. Third-mirror break-off volume snapshots are not supported.

You can:

- ◆ Create a storage pool (Creating a Storage Pool)
- ◆ Delete a storage pool (Deleting a Storage Pool)
- ◆ Add and delete storage pool disks (Adding and Removing Storage Pool Disks)
- ◆ Rename a storage pool (Renaming a Storage Pool)

Data Pool

The data pool is the first storage pool created within a disk group. All other storage pools that are subsequently created within a disk group are clone pools.

Note Only one data pool can be created within a disk group.

Clone Pools

A clone pool is a storage pool that contains one or more full-sized instant volume snapshot replicas of volumes within a data pool. (A volume snapshot is an image of a volume's contents at the moment in time that the snapshot was taken. See the *VERITAS Volume Manager Administrator's Guide* for more information about volume snapshots.)

Note A clone pool can be moved into another disk group but an individual snapshot volume cannot. If you want to move a clone pool into a different disk group in order to perform off-host processing, you must ensure that the clone pool contains only those snapshots that you require for this purpose.

Storage Pool Policy

A storage pool's policies define how it behaves when more storage is required, and when you try to create volumes whose capabilities are not permitted by the current templates.

Note A LUN, or *logical unit*, can either correspond to a single physical disk, or to a collection of disks that are exported as a single logical entity, or virtual disk, by a device driver or by an intelligent disk array's hardware. VEA and other VERITAS software modules may be capable of automatically discovering the special characteristics of LUNs, or you can use the Annotation Service to define new storage attributes.

The two policies that are associated with a storage pool are:

- ◆ `AutoGrow` Defines how a storage pool uses LUNs. This policy can take the following *level* values:
- ◆ `poolUse` storage that has been manually assigned to the storage pool.
- ◆ `diskgroupUse` storage that has been assigned to the disk group.
- ◆ `SelfSufficient` Defines how a storage pool uses templates. This policy can take the following *level* values:
- ◆ `poolUse` templates that have been manually assigned to the storage pool.

Managing Storage Pools

- ◆ `diskgroup` Use templates that have been manually assigned to the disk group.
- ◆ `host` Use templates that have been installed in the ISP repository on this host.

These values of these two policies can be combined to suit how the storage pool is to be used.

Storage Pool Set

A storage pool set is a bundled definition of the capabilities of a data pool and its clone pools. For convenience, you can use a storage pool set definition to define both pools in a single operation. See “Storage Pool Sets” on page 288 for details of the storage pools sets that are provided.

Creating a Storage Pool

Storage pools are associated with disk groups. Use the following procedure to create a storage pool using the templates provided with *ISP*.

Note If you do not want to use the provided templates to create your storage pool, then you need to create a disk group (see *Creating a Disk Group*) using the *Organize Diskgroup* option (refer also to *Organizing Disk Groups*).

▼ To create a storage pool

1. Select an *ISP* disk group.
2. Choose **Actions > New Storage Pool**. There is a short interval while the required data is downloaded from the server to launch the wizard.
3. When the **Storage Pool Wizard** appears, click **Next>**. The Storage Pool Information dialog box appears.
4. Complete the Storage Pool Information dialog box as follows:

Pool Name:	Enter the name for the storage pool. The name must be less than 31 characters.
DiskGroup:	Confirm the disk group, or select another disk group, to contain the storage pool.
Storage Pool Description:	If you want, you can enter a brief description for the new storage pool.

5. Click **Next>**. The Select Storage Pool dialog is displayed.
6. Select an existing storage pool or storage pool definitions to be used as a reference while creating the new storage pool. The new storage pool will have characteristics (policies, templates and capabilities) similar to the selected storage pool or storage pool definition. Storage pool definitions are listed at the bottom of storage pool list.
7. Click **Next>**. The Select Disks dialog is displayed.
8. The Select Disk dialog has two panes. The left pane (**Available Disks:**) shows the names and sizes of all available disks - those disks in the disk group that are not already part of any storage pool. Move the selected disks to the right pane (**Selected Disks:**).

Managing Storage Pools

9. Click **Next**. The summary page displays information you have entered. If you want to make changes, click **<Back**. Otherwise, click **Finish**.
10. If the storage pool is successfully created, the following message appear:
Storage Pool Created Successfully

Deleting a Storage Pool

Use the following procedure to delete an existing storage pool:

▼ To delete a storage pool

1. Select the *ISP* disk group containing the storage pool. The Disk Group view is displayed in the right pane.
2. Click on the Storage Pools tag. The storage pools are listed in the right pane.
3. Select the storage pool to be deleted.
4. Choose **Actions > Delete Storage Pool**. Click **Yes** when the Delete Storage Pool confirmation dialog box appears.

Adding and Removing Storage Pool Disks

Use the following procedure to add or remove one or more storage pool disks:

▼ To add or remove storage pool disks

1. Select the ISP disk group containing the storage pool. The Disk Group view is displayed in the right pane.
2. Click on the **Storage Pools** tag. The storage pools are listed in the right pane.
3. Select the storage pool containing the disks to be added or removed.
4. Choose **Actions > Add/Remove Disks**. There is a short interval while status information is downloaded. The **Add/Remove disks** wizard is displayed.
5. Select the **Disk Group** associated with the storage pool from the pulldown menu.
6. Select the **Pool Name** from which you want to remove disks from the pulldown menu.
7. The **Add/Remove disks** dialog appears. The **Add/Remove Disks** dialog has two panes. The left pane shows the names and sizes of all available disks. The right pane shows the selected disks.
8. Select the disks to be added or removed. To add disks to the pool, select the disks from the Available Disks and **Add>** them to the Selected Disk list. To remove disks from the pool, select the disks from the Selected Disks and **Remove<** them to the Available Disks pane.
9. Click **Next**. The summary page displays information you have entered. If you want to make changes, click **<Back**. Otherwise, click **Finish**.

Renaming a Storage Pool

Use the following procedure to rename an existing storage pool:

▼ **To rename a storage pool**

1. Select the *ISP* disk group containing the storage pool to be renamed. The Disk Group view is displayed in the right pane.
2. Click on the **Storage Pools** tag. The storage pools are listed in the right pane.
3. Select the storage pool to be renamed, then choose **Actions > Rename Storage Pool**. The Rename Storage Pool dialog appears.
4. Type in the new name; then click **OK**.

Associating a Template Set with Storage Pool

A template set is a logical collection of related templates. Associating a template set to a storage pool allows to change organization of storage pool. You can associate a template set with a storage pool.

▼ To associate a template set with a storage pool

1. Right click the storage pool, and select **Associate Template Set**. The Associate Template Set wizard appears.
2. Click **Next>**. The **Select template sets** dialog appears.
3. You can now change the storage pool, if you wish.
4. Select one or more template sets to associate with the storage pool by checking the appropriate checkboxes underneath **Template sets**.

Note For each template set you select to associate with a storage pool, a definition of that particular template set is located at the bottom of the **Select template sets** dialog box.

5. Click **Next>**.
6. The **Summary Page** displays the storage pool and the list of associated template sets.
7. Click **Finish** to complete the operation.

Disassociating a Template Set from a Storage Pool

A template set is a logical collection of related templates. Disassociating a template set to a storage pool allows to change organization of storage pool. You can disassociate a template set from a storage pool.

▼ To disassociate a template set from a storage pool

1. Right click the storage pool, and select **Disassociate Template Set**. The Disassociate Template Set wizard appears.
2. Click **Next>**. The **Select template sets** dialog appears.
3. You can now change the storage pool, if you wish.
4. Select one or more template sets to disassociate with the storage pool by checking the appropriate checkboxes underneath **Template sets**.

Note For each template set you select to disassociate with a storage pool, a definition of that particular template set is located at the bottom of the **Select template sets** dialog box.

5. Click **Next>**.
6. The **Summary Page** displays the storage pool and the list of disassociated template sets.
7. Click **Finish** to complete the operation.

Annotating a Disk

The **Annotate Disks** feature enables you to administer disk storage attributes.

Storage Attributes

A storage attribute allows the properties of a LUN to be defined in an arbitrary conceptual space. For example, attributes can describe properties such as:

- ◆ Manufacturer.
- ◆ Model type.
- ◆ Physical location, such as rack number, frame number, floor, building, or site.
- ◆ Hardware RAID configuration.
- ◆ Failover properties.
- ◆ Performance properties.

Attributes may be used to capture information about special features that storage possesses, such as:

- ◆ Hardware-supported cloning, such as EMC Business Continuity Volumes (BCV).
- ◆ Hardware-supported replication, such as the EMC Symmetrix Remote Data Facility (SRDF).
- ◆ Hardware redundancy, such as mirrored parity, and caching.

It should only be necessary to enter such information manually if VEA cannot discover it automatically. An example of a user-defined attribute is physical location.

Note Attribute names and their string values are case sensitive.

In order to administer storage attributes, you must first Launch the Annotate Service Dialog. You can then:

- ◆ Add a New User Attribute
- ◆ Modify a User Attribute
- ◆ Remove a User Attribute
- ◆ Sort Disks by Attribute Name
- ◆ Create a User Attribute Short-cut
- ◆ Remove a User Attribute

Note There are four *discovered* attributes: *name*, *enclosure*, *disk group* and *vendor*.

Launch the Annotate Service Dialog

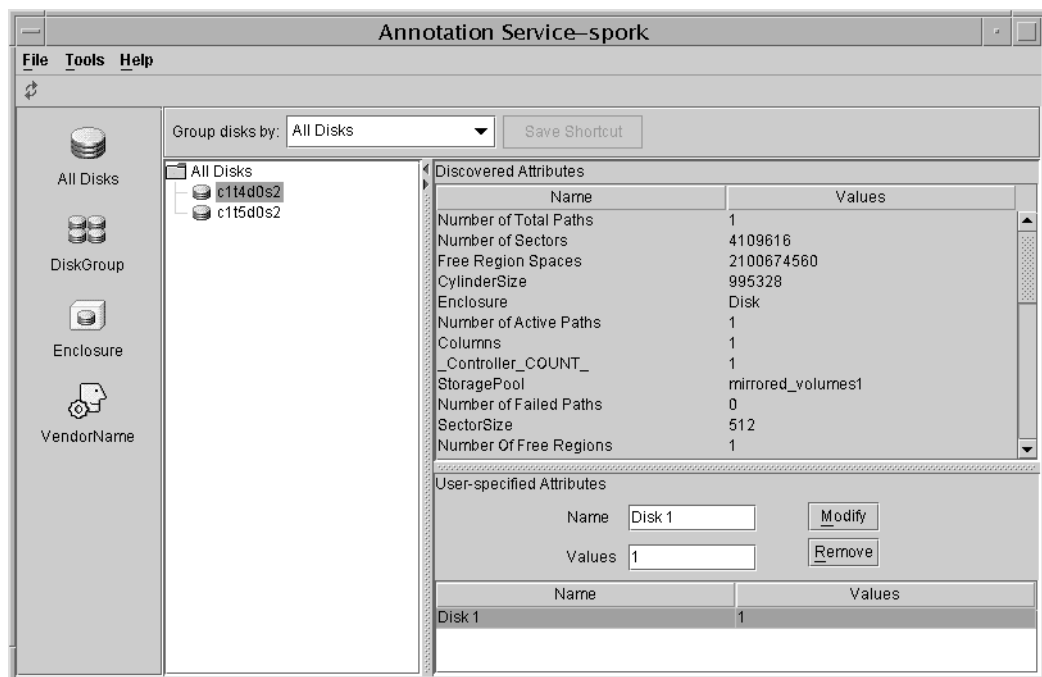
Use the following procedure to launch the Annotate Service dialog:

▼ To launch the Annotate Service Dialog

1. Select the **Disks** node.
2. Choose **Actions > Annotate Disks**. There is a short interval while status information is downloaded. The Annotate Service dialog appears.

The left pane lists all the disks available under the Volume Manager control. The right pane lists the name and values of both discovered attributes and user set attributes.

Annotate Service Dialog



3. Click on one of the shortcuts in the left pane to display the required disk set. The disks shown are grouped according to their attribute names and values represented by the shortcut information.

4. Choose one or more disks to which you want to annotate. (The existing attributes are shown in the attribute/value table.)

Note If you select multiple disks, only those attributes common to the selected disks are displayed.

Add a New User Attribute

Use the following procedure to add new user attributes:

▼ To add a new user attribute

1. Launch the Annotate Service Dialog and select the disks to be annotated.
2. Click on the **Add** button alongside the lower attribute/value table. A new row appears at the bottom of the attribute table.
3. Click on the **Name** cell, and enter the name of the attribute to be added.
4. Click on the **Value** cell, and enter a value for the attribute. (If you want to enter more than one value, separate the values with commas.)

Repeat step 2 to step 4 until all attributes have been added; or, click the **Add** button for the number of attributes you want to set.

5. Choose **Save** or **Save and Exit**.

Note The **Reset** button will discard the unsaved changes to the disk attributes and return the original values

Modify a User Attribute

Use the following procedure to modify user attributes:

▼ To modify a user attribute

1. Launch the Annotate Service Dialog and select the disks to be annotated from the box underneath the attribute selection fields.
2. Click on the **Value** cell to enter new values for the attribute. If you want to enter more than one value, separate the values by a comma.
3. Repeat step 2 until you have completed all the changes.
4. Click **Modify**.

Remove a User Attribute

You can remove only the user set attributes.

Use the following procedure to remove user attributes:

▼ To remove a user attribute

1. Launch the Annotate Service Dialog and select the disks to be annotated.
2. Click on the attributes to be removed.
3. Click on the **Remove** button alongside the lower attribute/value table.
4. Choose **Save** or **Save and Exit**.

Note Use the **Reset** button to discard unsaved changes and return the original values.
Use the **Cancel** button discard the changes and close the window.

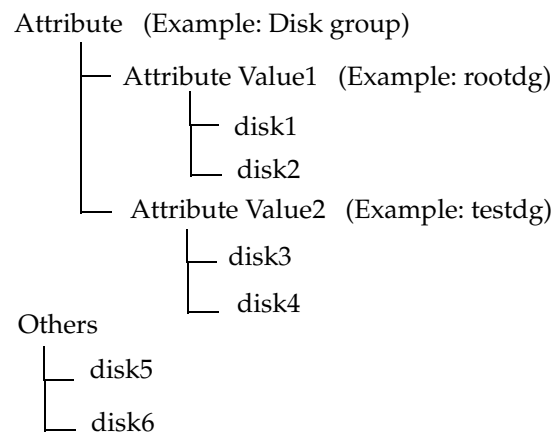
Sort Disks by Attribute Name

You can display the disks having the same attributes. You can either view using a short-cut (Create a User Attribute Short-cut), or by using the following procedure to sort by attribute name:

▼ To sort disks by attribute

1. Launch the Annotate Service Dialog.
2. Select the attribute from the combo box. The tree structure then changes to display the attribute followed by the attribute value and then the disks that have these values. Those disks that do not have the selected attribute, are listed under *Others*.

Attribute Tree



Create a User Attribute Short-cut

Short-cuts are specific to client machines. Short-cuts you create will stay in the annotation service window even after you disconnect from the server. If you connect to another server, these shortcuts will be effective. If the new server does not have a property for which the client machine has a short-cut, a message will be displayed.

User attribute short-cuts are created in the left pane. (Discovered attributes also have short-cuts here.) You can add the user set attributes or the discovered attributes to this list.

Use the following procedure to create a user attribute short-cut:

▼ To create a user attribute short-cut

1. Launch the Annotate Service Dialog.
2. From the drop-down list above the disk tree view, select the attribute you want to create a short-cut for. The tree displays a list of all disks that have the selected attribute.
3. Select **Tools>Save Shortcut**.

This creates a new shortcut icon in the left panel for the selected attribute.

Delete a User Attribute Short-cut

You can remove user-created short-cuts.

▼ To remove a user attribute short-cut

1. Launch the Annotate Service Dialog.
2. Click **Delete Shortcut**.

Cluster Tasks

11

VEA provides access to cluster-related features only if the optional cluster functionality is present.

The following tasks are associated with clusters:

- ◆ Working with Clusters
- ◆ Creating a Shared Disk Group
- ◆ Importing a Shared Disk Group
- ◆ Maintaining Primary Policies
- ◆ Setting the Primary Node

The Cluster Task Roadmap shows a summary of common cluster tasks.

Cluster Task Roadmap

The following table provides a summary of common VxVM cluster tasks.

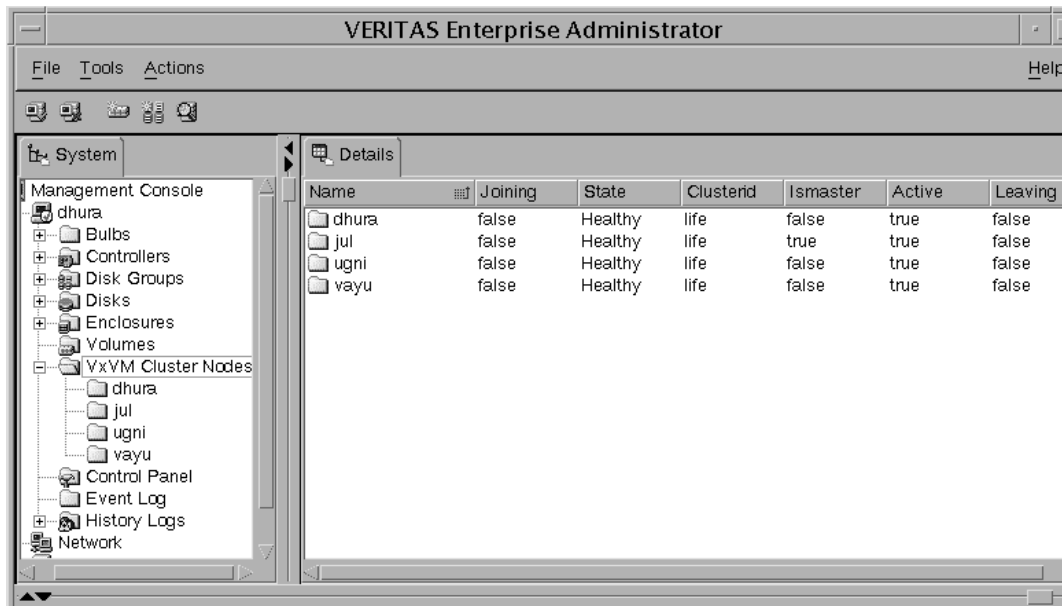
To perform this task ...	Use this procedure...
<i>Setup Tasks:</i>	
Setting the primary policy	Setting the Primary Policy
Clearing the primary policy	Clearing the Primary Policy
Setting the primary node	Setting the Primary Node
Set up a shared disk group	Creating a Shared Disk Group
<i>Other Tasks:</i>	
Access a deported disk group for cluster use	Importing a Shared Disk Group

Working with Clusters

VxVM includes an optional cluster feature that enables VxVM to be used in a cluster environment. To activate the cluster feature, you must obtain a license for the VxVM cluster functionality.

After adding the cluster license to the machine, you must restart the VxVM server to enable cluster support. When a cluster license is present on a machine, VEA displays a Clusters icon in the object tree (even if the machine is not currently set up as a node in a cluster). When a cluster is shown in the tree, all active/inactive nodes that are members of the cluster are displayed in the tree. This is illustrated in the figure, “Tree/Grid View of a Cluster.”

Tree/Grid View of a Cluster



- ◆ To display the objects in a cluster, expand the cluster hierarchy (by clicking on the plus sign next to the Clusters icon).
- ◆ To determine which node is the master in a cluster, expand the cluster hierarchy for the appropriate cluster and select **Hosts** from the cluster hierarchy. The grid displays the names and properties of the nodes in the cluster. The Type column indicates which node is the master.

The Cluster Properties window also indicates which node is the master. To access this window, select a cluster and choose **Properties** from the Selected or popup window.

Notes:

- ◆ When VEA is connected to a single host, that host is auto-expanded; that is, the hierarchy of objects under that host is fully displayed in the dialog.
- ◆ If VEA is connected to more than one host, the hierarchy of objects is not auto-expanded.

Creating a Shared Disk Group

The **New Disk Group** procedure creates an additional disk group. Disks must be placed in disk groups before they can be used by VxVM. You can create several disk groups to organize your disks into logical sets. For example, the disks in each disk group may have different performance or hardware characteristics, or they may be used for different purposes.

When you place a disk under VxVM control, the disk is initialized. Initialization destroys any existing data on the disk. (You can use encapsulation to preserve disk data.)

The steps for creating a new disk group are:

1. Right-click the Disk Groups folder, or right-click a disk.
2. Select **Actions > New Disk Group**. The New Disk Group wizard screen appears.
3. Click **Next** to continue. The New Disk Group dialog appears. Provide information as described in the following table.

Group name:	Enter a new disk group name in the Group Name field.
Create cluster group	If the cluster feature is available, you can select the Create cluster group checkbox if the new disk group is to be used with clusters.
Available disks: Selected disks:	Select which disks (from Available disks) you want to include in the group. Make sure the disks you want to include are in the right pane of the window (Selected disks). (Normally, you would add all the disks you want in the group at this point. However, you can always add more disks later with the Add Disk to Disk Group procedure.)
Disk names:	Type in the names of any disks to be added that do not appear under Selected disks .

Creating a Shared Disk Group

Create cluster group:	<input type="checkbox"/> Check this box to create a shared disk group.
-----------------------	------------------------------------------------------------------------

4. If your system supports Activation Mode, then select **Activation Mode** and choose the mode from the displayed submenu. Click **OK**.
5. The next screen confirms the disks you have selected. Choose **Yes** to continue if you are satisfied with the disk selection. If you are not satisfied, you can click the **No** button to go back to the previous screen in order to modify your disk choices.
6. Normally, you would add all the disks you want in the group at this point. You can always add more disks later with the **Add Disk to Disk Group** command.
7. If the disks are not empty, you now need to indicate whether the disk should be initialized or encapsulated. The initialize option is not available for disks with mounted partitions or swap areas. If you initialize the disk, any existing data on the disk will be destroyed. If you encapsulate the disk, any existing data will be preserved in volumes.
8. If you chose to encapsulate the disk, the disk needs to be rebooted before encapsulation can take effect. You have the option of rebooting the system now, or later.

Notes:

- ◆ The disk group name must be unique.
- ◆ The new disk group must contain at least one disk.
- ◆ Only disks that are online and do not belong to a disk group can be used to create a disk group.
- ◆ When a disk is placed under VxVM control, the disk is either initialized or encapsulated. Encapsulation preserves any existing data on the disk in the form of volumes. Initialization destroys any existing data on the disk. Encapsulation is recommended for the root/boot disk and any other disks that contain valuable data.

Importing a Shared Disk Group

The **Import Disk Group** procedure makes a deported (inaccessible) shared disk group and its volumes accessible again. To import a deported disk group, you must know the former name of the disk group, and this name must not be in use. In addition, at least one disk name formerly assigned to the deported disk group must not be in use.

▼ To import a shared disk group

1. Select the appropriate host machine node.
2. Choose **Actions > Import Disk Group**.
3. Complete the Disk Group dialog box as follows:

New Name:	Enter a new disk group name in the New Name field. to change the name of the disk group at import.
Clear host ID	Select Clear host ID to clear the existing host ID stamp on all disks in the disk group at import, Do not use this option if another host is using any disks in the disk group.
Force	Select Force to force the disk group import when the host cannot access all disks in the disk group This option can be used to import a disk group that contains a failed disk, but can lead to disk group inconsistency if all disks are still usable.
Import Shared	Select Import Shared to import the disk group as a shared disk group, This option is only applicable in a cluster environment.

Maintaining Primary Policies

When you have provided all necessary information in the dialog box, click **OK**.

Notes:

- ◆ Only deported disk groups can be imported.
- ◆ A deported disk group cannot be imported if another disk group with the same name has been created since the disk group was deported.
- ◆ When a disk group is imported, the system stamps its host ID on all disks in the disk group. A disk group import fails if one of the disks is stamped with a host ID that does not match the others. This ensures that dual-ported disks cannot be managed (and possibly corrupted) by two systems at the same time. If you are *sure* that the disk group is *not* in use by another host, you can clear the host IDs and import the disk group by selecting the **Clear Host ID** option.

Maintaining Primary Policies

You maintain primary policies by:

- ◆ Setting the Primary Policy
- ◆ Clearing the Primary Policy

Setting the Primary Policy

To set the primary policy for a file system, use the following procedure:

▼ **To set the primary policy**

1. Select the file system.
2. Choose **Actions > Cluster > Select Primary Policy**. The Set Primary Policy dialog is displayed.
3. Move the required primary nodes from the **All Cluster Nodes:** list to the **Primary Policy Nodes:** list.
4. Click **OK** to complete the setting.

Note This operation cannot be performed on clones or snapshots.

Clearing the Primary Policy

To clear the primary policy for a file system, use the following procedure:

▼ To clear the primary policy


1. Select the file system.
2. Choose **Actions** > **Cluster** > **Clear Primary Policy**. The confirmation dialog is displayed.
3. Click **OK**.

Setting the Primary Node



To set the primary node for a file system, use the following procedure:

▼ To set the primary node

1. Select the file system.
2. Choose **Actions** > **Cluster** > **Select Primary Node**. The Set Primary Node dialog is displayed.
3. Select the primary node from the drop-down menu.
4. Click **OK** to complete the setting.



Setting the Primary Node



Troubleshooting

12

Select one of the following options for information on alerts, object states, and other troubleshooting topics. If you are unable to resolve a problem by following the instructions provided, please contact Customer Support for additional information.

- ◆ Alerts
- ◆ Disk Troubleshooting
- ◆ Volume Troubleshooting
- ◆ General Troubleshooting
 - ◆ Client/Server Problems
 - ◆ Configuration Changes

Alerts

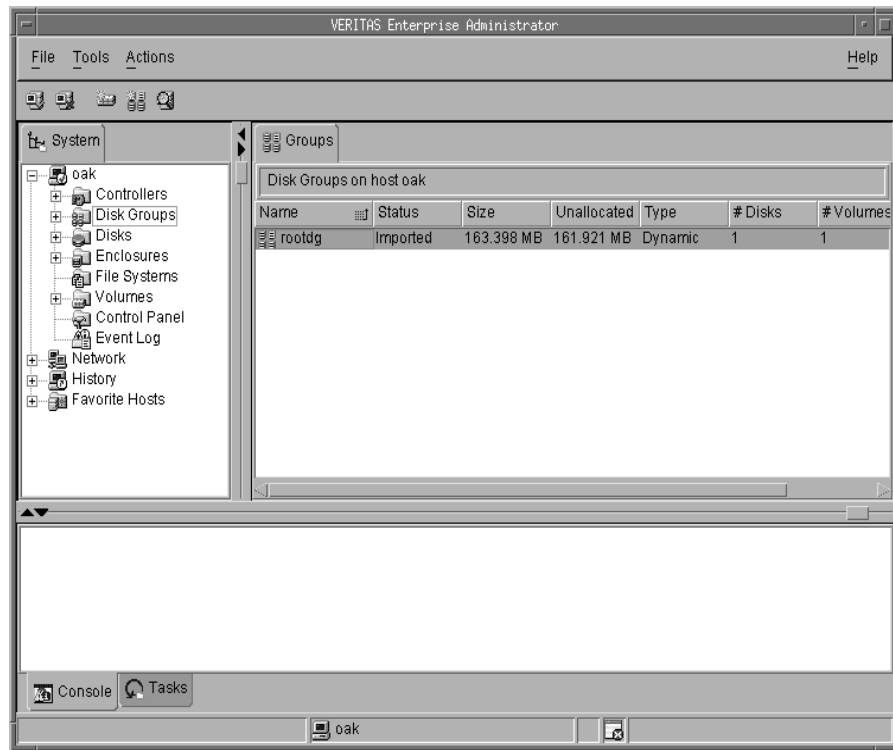
When an object fails or experiences an error, an alert icon appears on the object and the object state usually changes to an unhealthy state. When an object enters an unhealthy state and has an alert, you should review the contents of the console window and then try to correct the problem.

Client Java Exceptions and Errors

When a client Java exception or error occurs, an error alert icon is displayed on the bottom of the main window.

Client Java Exceptions and Errors

Error Alert Icon



Click on this icon to display the Java console containing the text of the errors or exceptions.

Disk Troubleshooting

The disk state indicates the current condition of a disk. The disk state appears in the state column of the grid and in the Disk Properties window. When a disk is unhealthy, an alert icon appears on the disk and the disk is listed in the Alert Monitor window. Disk states can be one of the following:

- ◆ Disk State: Imported
- ◆ Disk State: Free
- ◆ Disk State: Not Setup
- ◆ Disk State: Import Failed
- ◆ Disk State: Disconnected
- ◆ Disk State: Deported
- ◆ Disk State: Offline
- ◆ Disk State: Not Initialized
- ◆ Disk State: Locked
- ◆ Disk State: Inactive
- ◆ Disk State: Foreign Owned

Disk State: Imported

The disk is accessible and has no known problems. The disk is under VxVM control and belongs to a disk group. This is the normal disk state. No user action is required.

Disk State: Free

The disk is freely available. A disk that is freely available is under VxVM control, but does not belong to a disk group.

- ◆ To make the disk available for VxVM use, add the disk to a disk group.

Disk State: Not Setup

The disk is not under VxVM control.

- ◆ To make the disk available for VxVM use, place the disk under VxVM control and add the disk to a disk group.

Disk State: Import Failed

The disk was previously under VxVM control, but VxVM cannot access the disk and/or configuration information for the disk. The disk is not available for use. This can happen after you upgrade VxVM. This can also happen if some disks are inaccessible at system startup.

- ◆ Make sure that the physical disk is turned on, plugged in, and attached to the computer. Repair any disk, controller, or cable problems. Scan the disks on the system (**Actions > Rescan**) to make sure that VxVM recognizes all attached hardware. Try to import the disk group for the disk.
- ◆ If the disk group is imported and the disk does not have a VxVM disk name, place the disk under VxVM control and add the disk to a disk group.

Disk State: Disconnected

The connection between the VxVM disk and the underlying physical disk (device) is broken. This can occur when a disk fails or when you disconnect a disk.

- ◆ Make sure that the physical disk is turned on, plugged in, and attached to the computer. Repair any disk, controller, or cable problems. Scan the disks on the system (**Actions > Rescan**) to make sure that VxVM recognizes any newly attached hardware.
- ◆ Try to reconnect the disconnected VxVM disk to its physical disk (using the **Replace Disk** task).
- ◆ If you cannot fix the problem, replace the disk with another physical disk.

Disk State: Deported

The disk belongs to a disk group that is currently deported. The disk is not available for use in the deported disk group and cannot be added to a new disk group.

- ◆ To return the disk to a usable state, import the disk group for the disk.
- ◆ To move the disk to another disk group, import the current disk group, remove the disk, and add the disk to a new disk group.
- ◆ If the disk group import fails and you want to free the disk space for use in other disk groups, you can use the **Initialize Disk** task.

Caution This command overwrites the disk header and configuration information on the disk. Any volume copies or regions located on the disk are not accessible after you reinitialize the disk.

Disk State: Offline

VxVM cannot access the disk.

- ◆ To return the disk to an accessible state, bring the disk online.

Disk State: Not Initialized

A configuration error prevents the disk from being used.

- ◆ To make the disk available for VxVM use, place the disk under VxVM control and add the disk to a disk group.
- ◆ If the disk cannot be added to VxVM, it can be necessary to reinitialize the disk with the **Initialize Disk** task.

Caution This command overwrites the disk header and configuration information on the disk. Any volume copies or regions located on the disk are not accessible after you reinitialize the disk.

Disk State: Locked

The disk belongs to a disk group that is currently imported by another machine or was last imported by another machine and was not deported cleanly. The disk is not currently usable on the local machine.

- ◆ To make all of the disks in the disk group available for use on the local machine, deport the disk group on the remote machine and then import the disk group on the local machine.
- ◆ To make one disk available for use on the local machine, remove the disk from its disk group (on the other machine) and add it to a disk group on the local machine.
- ◆ If the disk group is *not* imported on another machine, it can be necessary to re-import the disk group and clear the host ID on all disks in the disk group. You can do this by selecting the **Clear Host ID** option in the Import Disk Group dialog box.

Caution If any of the disks in the disk group are in use by another machine, this command can cause the disks to be managed by two systems at the same time, which can lead to corruption.

Disk State: Inactive

The disk appears to belong to a disk group, but is not active in the disk group. The disk may have become (temporarily) inaccessible. Alternatively, the disk may have been inaccessible at disk group import, but may be accessible now.

- ◆ Make sure that the physical disk is turned on, plugged in, and attached to the computer. Repair any disk, controller, or cable problems. Scan the disks on the system (**Rescan**) to make sure that VxVM recognizes any newly attached hardware.
- ◆ It may be necessary to reattach the disk to its disk group with the **Replace Disk** task.
- ◆ After a successful reattach, volume recovery can be necessary.
- ◆ If the reattach fails and you want to free the disk space for use, you can use the **Initialize Disk** task.

Caution This command overwrites the disk header and configuration information on the disk. Any volume copies or regions located on the disk is not accessible after you reinitialize the disk.

Disk State: Foreign Owned

The disk is under the control of another disk management product. The disk is not available for use with VxVM.

- ◆ To make the disk available for VxVM use, remove the disk from the other disk management product control and place the disk under VxVM control.

Volume Troubleshooting

The volume state indicates the current condition of a volume. The volume state appears in the state column of the grid and in the Volume Properties window. The normal state is "Healthy," which means that the volume is accessible and has no known problems.

When a volume is unhealthy, an alert icon appears on the volume and the volume is listed in the Alert Monitor window. Unhealthy states include:

- ◆ Volume State: Stopped
- ◆ Volume State: Uninitialized
- ◆ Volume State: Degraded
- ◆ Volume State: Failed

Volume State: Stopped

The volume is disabled and is not available for use.

- ◆ To return the volume to a usable state, start the volume.

Volume State: Uninitialized

The volume is enabled, but is not initialized. If this state appears during volume creation, no user action is necessary.

- ◆ To put the volume in a usable state, start the volume.

Volume State: Degraded

The RAID-5 volume is in a degraded mode. This usually occurs when one or more subdisks are unavailable and read requests require data reconstruction. An underlying disk may have failed completely or there may be I/O errors on part of a disk. Alternatively, an underlying disk may be disconnected.

- ◆ Make sure that the physical disks are turned on, plugged in, and attached to the computer. Repair any disk, controller, or cable problems. Scan the disks on the system (**Rescan**) to make sure that VxVM recognizes any newly attached hardware.
- ◆ To return the RAID-5 volume to a healthy state, recover the volume.
- ◆ If there has been a partial or complete disk failure, either replace the disk or move the affected subdisks to another disk.

Volume State: Failed

The RAID-5 volume is unusable. This usually occurs when there is a double disk failure (that is, when two subdisks in the same stripe or the parity and one subdisk in a stripe are damaged or inaccessible). The underlying disks may have failed or become inaccessible.

- ◆ Make sure that the physical disks are turned on, plugged in, and attached to the computer. Repair any disk, controller, or cable problems. Scan the disks on the system (**Rescan**) to ensure that VxVM recognizes any newly attached hardware.

Run the following command to reattach the disks to their disk group:

```
# /etc/vx/bin/vxreattach [device_name]
```

To return the RAID-5 volume to a healthy state, recover the volume.

- ◆ If two of the volume disks have failed, the RAID-5 volume data is unusable. Replace the failed disks and then try to restart the volume by running the following command:

```
# vxvol -f start volume_name
```

Restore the volume from backup.

To prevent this problem in the future, enable logging for RAID-5 volumes.

General Troubleshooting

General troubleshooting includes:

- ◆ Client/Server Problems
- ◆ Configuration Changes

Client/Server Problems

User is not authorized to connect to the server.

This error can occur for the following reasons:

- ◆ The user name is incorrect.
- ◆ The password is incorrect.
- ◆ The user does not have an account on the machine.
- ◆ The password and group files or NIS tables do not contain entries for the user.
- ◆ The user is not listed in the `vrtsadm` group entry.

Enter the user name and password again. If the same error appears, ask the administrator of the machine to make sure that your user name is included in the `vrtsadm` group entry and your user name and password are entered correctly in the appropriate password and group files or NIS tables.

There is no such server.

The specified host machine does not exist or is not in the Domain Name System (DNS).

- ◆ Check the machine name and enter the machine name again.

Unable to connect to the server.

The VEA system may not be running on the machine to be administered.

- ◆ Make sure the server is running and the client and server are properly configured. If necessary, restart the VEA system.

Possible client/server versions mismatch.

The VEA client and server versions may be incompatible.

- ◆ Check the client and server versions using the following commands:

```
# vxsvc -V
```

If necessary, upgrade the client or server.

Server connection lost.

The VEA service has been stopped.

- ◆ Restart the server.
- ◆ Use the following command to determine if the server connection is lost.

```
# vxsvc
```

The server seems to have multiple addresses.

The server (host) may be known by multiple addresses.

- ◆ Try to specify one of the server IP addresses as the server host name.

VEA version of plug-in component is not compatible. Cannot load plug-in

Trying to connect from a pre 4.0 version of the VEA Console (such as, from a 3.5 client to a 4.0 server).


- ◆ Upgrade to a later VEA client.

Configuration Changes



VxVM is dynamic, so it automatically displays new objects or changes to existing objects on the system. These changes or additions are usually reflected in VEA immediately (or within a short time).

If you add objects (such as new disks) to the system and VxVM does not display these objects within a short time, you can force VxVM to check for new objects in one of the following ways:

1. Scan the disks on the system (**Rescan Disks**) to make sure that VxVM recognizes any newly attached hardware.
2. Update the system configuration database.



General Troubleshooting






Storage Selection Rules

A

The following table lists the rules that are presented for storage selection in ISP.

Rules for Storage Selection	Meaning
Don't select the storage satisfying any of:	Prevents the selected set of LUNs from participating in the allocation rules.
Select the storage satisfying all of:	Confines the storage to all of the specified set of LUNs.
Select the storage satisfying any one of:	Confines the storage to any one of the specified set of LUNs.
Select the storage satisfying one or more of:	Confines the storage to any (one or more) of the specified set of LUNs.
Select the storage that is separated by:	Confines the storage by the attributes specified. This rule does not accept attribute values. It does not support LUNs as operands. Only attribute name is supported.



Available ISP Definitions

B

This appendix lists and describes the following types of pre-defined objects that are available for use with VERITAS Intelligent Storage Provisioning (ISP):

- ◆ Template Sets
- ◆ Volume Templates
- ◆ Capabilities
- ◆ Storage Pools
- ◆ Storage Pool Sets

These objects are defined in the file `/etc/vx/alloc/alloc_capabilities.txt`. Before modifying this file, make a backup copy so that you can reverse any changes that you make.

Template Sets

ConfineVolume

Ensures that a volume is confined to specified storage.

Provides capabilities: ConfineToSimilarStorage, ConfineToSpecificStorage.

Uses templates: ConfineToSimilarStorage, ConfineToSpecificStorage.

DataMirroring

Ensures that a volume has multiple copies of its data.

Provides capabilities: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents.

Uses templates: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents.

DataMirroringPrefabricatedRaid5

Ensures that a volume has multiple copies of data on prefabricated RAID-5 disks that are exported by an array.

Provides capabilities: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, PrefabricatedRaid5.

Uses templates: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, PrefabricatedRaid5.

DataMirroringPrefabricatedStriping

Ensures that a volume has multiple copies of data on prefabricated striped disks that are exported by an array.

Provides capabilities: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, PrefabricatedStriping.

Uses templates: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, PrefabricatedStriping.

DataMirrorStripe

Ensures that I/O from and to a volume is spread across multiple columns within mirrors.

Provides capabilities: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataMirrorStripe, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents.

Uses templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataMirrorStripe, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents.

DataStripeMirror

Ensures that I/O from and to a volume is spread across multiple columns, where each column has multiple copies of the data.

Provides capabilities: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataStripeMirror, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents.

Uses templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, DataStripeMirror, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents.

InstantSnapshottable

Ensures that a volume supports dirty region logging (DRL) and instant snapshots.

Provides capabilities: ConfineLogsToSimilarStorage, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents.

Template Sets

Uses templates: `ConfineLogsToSimilarStorage`, `DCOLogMirroring`, `InstantSnapshottable`, `LogsOnSeparateComponents`.

MultipathingThroughMirroring

Ensures that a volume can withstand the failure of a number of paths. The I/O from and to a volume can potentially be spread across all the paths.

Provides capability: `DataMirroring`, `Multipathing`.

Uses template: `MultipathingThroughMirroring`.

MultipathingThroughMultiplePaths

Ensures that a volume can withstand the failure of a specified number of paths. The I/O from and to the volume can potentially be spread across all these paths.

Provides capability: `MultipathingThroughMultiplePaths`.

Uses template: `MultipathingThroughMultiplePaths`.

PrefabricatedDataMirroring

Ensures that a volume uses prefabricated data mirroring configured within the disks exported by an array.

Provides capabilities: `ArrayProductId`, `PrefabricatedDataMirroring`.

Uses templates: `ArrayProductId`, `PrefabricatedDataMirroring`.

PrefabricatedRaid5

Ensures that a volume uses prefabricated RAID-5 disks that are exported by an array.

Provides capabilities: `ArrayProductId`, `PrefabricatedRaid5`.

Uses templates: `ArrayProductId`, `PrefabricatedRaid5`.

PrefabricatedStriping

Ensures that a volume uses prefabricated striped disks that are exported by an array.

Provides capabilities: `ArrayProductId`, `PrefabricatedStriping`.

Uses templates: `ArrayProductId`, `PrefabricatedStriping`.

Raid5Templates

Ensures that a volume uses parity to maintain redundant data.

Provides capabilities: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, LogsOnSeparateComponents, Raid5Capability, Raid5LogMirroring.

Uses templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, LogsOnSeparateComponents, Raid5Volume.

Striping

Ensures that I/O from and to volume is spread across multiple columns.

Provides capabilities: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, Striping.

Uses templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, Striping.

StripingPrefabricatedDataMirroring

Ensures that I/O from and to a volume are spread across multiple columns, which are configured on prefabricated data mirror disks that are exported by an array.

Provides capabilities: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, PrefabricatedDataMirroring, Striping.

Uses templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, PrefabricatedDataMirroring, Striping.

Volume Templates

ArrayProductId

Provides capability: ArrayProductId.

ColumnsOnSeparateComponents

Provides capability: ColumnsOnSeparateComponents.

ConcatVolumes

Provides Capability: ConcatVolumes

ConfineColumnsToSimilarStorage

Provides capability: ConfineColumnsToSimilarStorage.

ConfineLogsToSimilarStorage

Provides capability: ConfineLogsToSimilarStorage.

ConfineMirrorsToSimilarStorage

Provides capability: ConfineMirrorsToSimilarStorage.

ConfineToSimilarStorage

Provides capability: ConfineToSimilarStorage.

ConfineToSpecificStorage

Provides capability: ConfineToSpecificStorage.

DataMirroring

Provides capability: DataMirroring.

Variable:nmirrsNumber of mirrors.

DataMirrorStripe

Provides capability: DataMirrorStripe.

Variables:ncolsMinimum number of columns.

nmaxcolsMaximum number of columns.

nmirrsNumber of mirrors.

DataStripeMirror

Provides capability: DataStripeMirror.

Variables:ncolsMinimum number of columns.

nmaxcolsMaximum number of columns.

nmirrsNumber of mirrors.

DCOLogMirroring

Provides capability: DCOLogMirroring.

DCOLogStriping

Provides capability: DCOLogStriping.

ExcludeSpecificStorage

Provides capability: ExcludeSpecificStorage

InstantSnapshottable

Provides capability: InstantSnapshottable.

LogsOnSeparateComponents

Provides capability: LogsOnSeparateComponents.

MirrorsOnSeparateComponents

Provides capability: MirrorsOnSeparateComponents.

MultipathingThroughMirroring

Provides capabilities: DataMirroring, Multipathing.

Variable:nmirrorsNumber of mirrors.

MultipathingThroughMultiplePaths

Provides capability: MultipathingThroughMultiplePaths.

PrefabricatedDataMirroring

Provides capability: PrefabricatedDataMirroring.

PrefabricatedRaid5

Provides capability: PrefabricatedRaid5.

PrefabricatedStriping

Provides capability: PrefabricatedStriping.

Raid5LogStriping

Provides capability: Raid5LogStriping.

Raid5Volume

Provides capabilities: Raid5Capability, Raid5LogMirroring.

Variables:ncolsMinimum number of columns.

nlogsNumber of logs.

nmaxcolsMaximum number of columns.

Striping

Provides capability: Striping.

Variables:`ncols`Minimum number of columns.

`maxcols`Maximum number of columns.



Capabilities

ArrayProductId

A volume uses storage of the same type (`productId`).

Provided by volume template: `ArrayProductId`.

ColumnsOnSeparateComponents

The columns of a volume are separated at the specified component level, such as "Controller" or "Enclosure".

Provided by volume template: `ColumnsOnSeparateComponents`.

Variable: `componentName` of component (default value is "Controller").

ConcatVolumes

A concatenated volume is created.

Provided by volume template: `ConcatVolumes`.

ConfineColumnsToSimilarStorage

Each column uses only storage that has the same value for a specified attribute.

Provided by volume template: `ConfineColumnsToSimilarStorage`.

Variable: `nameName` of storage attribute (default value is "Enclosure").

ConfineLogsToSimilarStorage

Each log uses only storage that has the same value for a specified attribute.

Provided by volume template: `ConfineLogsToSimilarStorage`.

Variable: `nameName` of storage attribute (default value is "Enclosure").

ConfineMirrorsToSimilarStorage

Each mirror uses only storage that has the same value for a specified attribute.

Provided by volume template: `ConfineMirrorsToSimilarStorage`.

Variable:name Name of storage attribute (default value is "Enclosure").

ConfineToSimilarStorage

A volume uses only storage that has the same value for a specified attribute.

Provided by volume template: ConfineToSimilarStorage.

Variable:name Name of storage attribute (default value is "VendorName").

ConfineToSpecificStorage

A volume uses only storage that has the specified value for a specified attribute name.

Provided by volume template: ConfineToSpecificStorage.

Variables:name Name of storage attribute (no default).

value Value of storage attribute (no default).

DataMirroring

A mirrored volume maintains multiple copies of its data. This capability extends the DataRedundancy capability.

Provided by volume template: DataMirroring.

Variable:nmirrs Number of mirrors (default value is 2).

DataMirrorStripe

A mirrored-stripe volume distributes I/O across multiple columns within mirrors. This capability extends the DataMirroring and Striping capabilities.

Provided by volume template: DataMirrorStripe.

Variables:ncols Minimum number of columns (default value is 8).

nmaxcols Maximum number of columns (default value is 20).

nmirrs Number of mirrors (default value is 2).

DataRedundancy

A volume that maintains redundant data.

Capabilities

Provided by capability: Raid5Capability.

Extended by capability: DataMirroring.

DataStripeMirror

A striped-mirror volume distributes I/O across multiple columns, where each column has multiple copies of data. This capability extends the DataMirroring and Striping capabilities.

Provided by volume template: DataStripeMirror.

Variables: `ncols` Minimum number of columns (default value is 8).

`nmaxcols` Maximum number of columns (default value is 20).

`nmirs` Number of mirrors (default value is 2).

DCOLogMirroring

The number of DCO plexes (copies) to configure for a DCO volume.

Provided by volume template: DCOLogMirroring.

Variable: `nlogs` Number of DCO plexes (default value is 1).

DCOLogStriping

The number of columns to configure for a striped DCO volume.

Provided by volume template: DCOLogStriping.

Variables: `ndccols` Minimum number of columns (default value is 4).

`nmaxdccols` Maximum number of columns (default value is 10).

ExcludeSpecificStorage

A volume does not use storage that has the specified value for the specified attribute name.

Provided by volume template: ExcludeSpecificStorage.

Variable: `name` Name of the storage attribute (no default).

Variable: `value` Value of the storage attribute (no default).

InstantSnapshottable

A volume supports instant snapshots (full-sized or space-optimized). This capability extends the Snapshottable capability.

Provided by volume template: InstantSnapshottable.

LogsOnSeparateComponents

The logs of a volume are separated at the specified component level, such as "Controller" or "Enclosure".

Provided by volume template: LogsOnSeparateComponents.

Variable:componentName of component (default value is "Enclosure").

MirrorsOnSeparateComponents

The mirrors of a volume are separated at the specified component level, such as "Controller" or "Enclosure".

Provided by volume template: MirrorsOnSeparateComponents.

Variable:componentName of component (default value is "Enclosure").

Multipathing

Multipathing allows a volume to withstand failure of a number of paths to the disks. The I/O from and to the volume can potentially be spread across all the paths.

Extended by capability: MultipathingThroughMultiplePaths.

MultipathingThroughMultiplePaths

Multipathing allows a volume to withstand failure of the specified number of paths to the disks. The I/O from and to the volume can potentially be spread across all the paths. This capability extends the Multipathing capability.

Provided by volume template: MultipathingThroughMultiplePaths.

Variable:npathsNumber of paths that are allowed to fail (default value is 2).

PrefabricatedDataMirroring

A volume uses prefabricated data mirroring that is configured on disks that are exported by an array. This capability extends the PrefabricatedDataRedundancy capability.

Provided by volume template: PrefabricatedDataMirroring.

Variable:`nmirrors`Number of prefabricated mirrors to use (default value is 2).

PrefabricatedDataRedundancy

A volume uses prefabricated redundant disks that are exported by an array.

Extended by capabilities: PrefabricatedDataMirroring, PrefabricatedRaid5.

PrefabricatedRaid5

A volume uses prefabricated RAID-5 disks that are exported by an array. This capability extends the PrefabricatedDataRedundancy capability.

Provided by volume template: PrefabricatedRaid5.

PrefabricatedStriping

A volume uses prefabricated striped disks that are exported by an array.

Provided by volume template: PrefabricatedStriping.

Raid5Capability

A RAID-5 volume uses parity to implement data redundancy. This capability provides the DataRedundancy capability.

Provided by volume template: Raid5Volume.

Variable:`ncolumns`Minimum number of columns (default value is 8).

Variable:`nmaxcolumns`Maximum number of columns (default value is 20).

Raid5LogMirroring

The number of RAID-5 log copies to configure for a RAID-5 volume.

Provided by volume template: Raid5Volume.

Variable:`nlogs`Number of RAID-5 log copies (default value is 1).

Raid5LogStriping

The number of columns to configure for striped RAID-5 logs.

Provided by volume template: Raid5LogStriping.

Variables: `nraid5cols` Minimum number of columns (default value is 4).
`nmaxraid5cols` Maximum number of columns (default value is 10).

Snapshottable

Volume snapshots of a volume may be taken.

Extended by capability: InstantSnapshottable.

Striping

A striped volume distributes I/O across multiple columns.

Provided by volume template: Striping.

Variables: `ncols` Minimum number of columns (default value is 8).
`nmaxcols` Maximum number of columns (default value is 20).

Storage Pools

any_volume_type

Supports volumes using any template from the configuration database.

Default policies: autogrow=2
 selfsufficient=3

Associate volume template: none

mirror_stripe_volumes

Supports distribution of I/O from and to volumes across multiple columns within mirrors.

Default policies:autogrow=2
 selfsufficient=1

Associates volume templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, ConfineToSimilarStorage, ConfineToSpecificStorage, DataMirrorStripe, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, MultipathingThroughMultiplePaths.

mirrored_prefab_raid5_volumes

Supports use by volumes of multiple copies of data configured on prefabricated Raid-5 disks that are exported by an array.

Default policies:autogrow=2
 selfsufficient=1

Associates volume templates: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, ConfineToSimilarStorage, ConfineToSpecificStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, MultipathingThroughMultiplePaths, PrefabricatedRaid5.

mirrored_prefab_striped_volumes

Supports use by volumes of multiple copies of data configured on prefabricated striped disks that are exported by an array.

Default policies:autogrow=2

```
selfsufficient=1
```

Associates volume templates: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, ConfineToSimilarStorage, ConfineToSpecificStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, MultipathingThroughMultiplePaths, PrefabricatedStriping.

mirrored_volumes

Supports volume with multiple copies of data.

Default policies:autogrow=2

```
selfsufficient=1
```

Associates volume templates: ArrayProductId, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, ConfineToSimilarStorage, ConfineToSpecificStorage, DataMirroring, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, MultipathingThroughMultiplePaths.

prefab_mirrored_volumes

Supports use by volumes of prefabricated data mirrors that are exported by an array.

Default policies:autogrow=2

```
selfsufficient=1
```

Associates volume templates: ArrayProductId, ConfineToSimilarStorage, ConfineToSpecificStorage, InstantSnapshottable, MultipathingThroughMultiplePaths, PrefabricatedDataMirroring.

prefab_raid5_volumes

Supports use by volumes of prefabricated RAID-5 disks that are exported by an array.

Default policies:autogrow=2

```
selfsufficient=1
```

Storage Pools

Associates volume templates: ArrayProductId, ConfineToSimilarStorage, ConfineToSpecificStorage, InstantSnapshottable, MultipathingThroughMultiplePaths, PrefabricatedRaid5.

prefab_striped_volumes

Supports use by volumes of prefabricated striped disks that are exported by an array.

Default policies:autogrow=2

```
selfsufficient=1
```

Associates volume templates: ArrayProductId, ConfineToSimilarStorage, ConfineToSpecificStorage, InstantSnapshottable, MultipathingThroughMultiplePaths, PrefabricatedStriping.

raid5_volumes

Supports volumes which use parity to maintain redundant data.

Default policies:autogrow=2

```
selfsufficient=1
```

Associates volume templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, ConfineToSimilarStorage, ConfineToSpecificStorage, InstantSnapshottable, LogsOnSeparateComponents, Raid5Volume, MultipathingThroughMultiplePaths.

stripe_mirror_volumes

Supports distribution of I/O from and to volumes across multiple columns, where each column has multiple copies of data.

Default policies:autogrow=2

```
selfsufficient=1
```

Associates volume templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineLogsToSimilarStorage, ConfineMirrorsToSimilarStorage, ConfineToSimilarStorage, ConfineToSpecificStorage, DataStripeMirror, DCOLogMirroring, InstantSnapshottable, LogsOnSeparateComponents, MirrorsOnSeparateComponents, MultipathingThroughMultiplePaths.

striped_prefab_mirrored_volumes

Supports distribution of I/O from and to volumes across multiple columns that are configured on prefabricated data mirrors exported by an array.

Default policies:autogrow=2

selfsufficient=1

Associates volume templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineToSimilarStorage, ConfineToSpecificStorage, InstantSnapshottable, MultipathingThroughMultiplePaths, PrefabricatedDataMirroring, Striping.

striped_volumes

Supports distribution of I/O from and to volumes across multiple columns.

Default policies:autogrow=2

selfsufficient=1

Associates volume templates: ArrayProductId, ColumnsOnSeparateComponents, ConfineColumnsToSimilarStorage, ConfineToSimilarStorage, ConfineToSpecificStorage, InstantSnapshottable, MultipathingThroughMultiplePaths, Striping.

Storage Pool Sets

mirrored_data_striped_clones

Supports data volumes with multiple copies of data, and snapshot volumes with I/O distributed across multiple columns.

Data storage pool type: mirrored_volumes.

Clone storage pool type: striped_volumes.

mirrored_prefab_raid5_data_mirrored_clones

Supports data volumes with multiple copies of data created from RAID-5 storage, and snapshot volumes with multiple copies of data.

Data storage pool type: mirrored_prefab_raid5_volumes.

Clone storage pool type: mirrored_volumes.

mirrored_prefab_stripe_data_striped_clones

Supports data volumes with multiple copies of data configured on prefabricated striped storage, and snapshot volumes with I/O distributed across multiple columns.

Data storage pool type: mirrored_prefab_striped_volumes.

Clone storage pool type: striped_volumes.

prefab_mirrored_data_prefab_striped_clones

Supports data volumes with multiple copies of data configured on prefabricated striped storage, and snapshot volumes with I/O distributed across multiple columns also configured on prefabricated striped storage.

Data storage pool type: prefab_mirrored_volumes.

Clone storage pool type: prefab_striped_volumes.

stripe_mirrored_data_striped_clones

Supports I/O from and to data volumes that are distributed across multiple columns, where each column has multiple copies of data, and snapshot volumes with I/O distributed across multiple columns.

Data storage pool type: stripe_mirror_volumes.


Clone storage pool type: striped_volumes.

striped_prefab_mirrored_data_striped_clones



Supports I/O from and to data volumes distributed across multiple columns constructed from mirrored storage, and snapshot volumes with I/O distributed across multiple columns.

Data storage pool type: striped_prefab_mirrored_volumes.

Clone storage pool type: striped_volumes.



Storage Pool Sets



Glossary

Action menu

A context-sensitive menu that changes its menu options to match the type of object that is selected. By default, the **Action** menu is greyed out. When an object is selected, the **Action** menu provides access to tasks appropriate for the selected object.

alert

An indication that an error or failure has occurred on an object on the system. When an object fails or experiences an error, an alert icon appears.

alert icon

An icon that indicates that an error or failure has occurred on an object on the system. Alert icons usually appear in the status area of the main window and on the group icon of the affected object.

Alert Monitor

A window that provides information about objects that have failed or experienced errors.

Array Support Library

See ASL.

ASL

A dynamically loadable library to discover attributes of a disk array.

browse dialog box

A dialog box that is used to view and/or select existing objects on the system. Most browse dialog boxes consist of a tree and grid.

button

A window control that the user clicks to initiate a task or display another object (such as a window or menu).

capability

You can create volumes having specific capabilities. A capability is a simple name used to indicate what the capability means. Storage is allocated according to your request. For example, you can ask for a volume to be Performant and Reliable. These are capabilities have parameters such as, number of striped columns or number of mirrors.

CFS

Cluster file system. A VxFS file system mounted on a selected volume in cluster (*shared*) mode.

check box

A control button used to select optional settings. A check mark usually indicates that a check box is selected.

children

Objects that belong to an object group.

cluster

A set of host machines (*nodes*) that shares a set of disks.

cluster file system

See CFS.

command log

A log file that contains a history of VERITAS Volume Manager (VxVM) tasks performed in the current session and previous sessions. Each task is listed with the task originator, the start/finish times, the task status, and the low-level commands used to perform the task.

DCO Data Change Object

See DCO.

DCO

A Volume Manager object that keeps the Persistent FastResync tracking information

DDI

DDI properties are device driver properties primarily used in Solaris OS.

DDL

Device Discovery Layer. A facility of VxVM for discovering disk attributes needed for VxVM DMP operation.

Device Discovery Layer

See DDL

Dirty Region Logging

See DRL.

device name

The physical disk device name (or *disk access name*).

dialog box

A window in which the user submits information to VxVM. Dialog boxes can contain selectable buttons and/or fields that accept information.

disk group

A set of disks that are under VxVM control and share a common configuration. A disk group configuration is a set of records containing detailed information on existing VERITAS Volume Manager objects (such as disk and volume attributes) and their relationships. Each disk group has an administrator-assigned name. Volumes can only be created on disks that belong to disk groups.

disk name

The VERITAS Volume Manager disk name (or *disk media name*). This is the logical name given to the disk by VxVM or the administrator.

DMP

The VERITAS Volume Manager Dynamic Multipathing facility.

dock

To separate or attach the main window and a subwindow.

DRL

A VxVM mechanism that tracks mirror write inconsistency which helps in fast resynchronization of mirrors after system crash. Even though the term *log* is used, internally DRL is implemented as a *map*.

Dynamic Multipathing

See DMP.

enclosure

A disk array.

FastResync

A VxVM feature for fast resynchronizing of stale mirrors at the time of reattach which were earlier split off as a result of snapshot or detach operation.

gap

A disk region that does not contain VERITAS Volume Manager objects (subdisks).

graphical view

A window that displays a graphical view of objects. In VxVM, the graphical views include the Object View window and the Volume Layout Details window.

grid

A tabular display of objects and their properties. The grid lists VERITAS Volume Manager objects, disks, controllers, or file systems. The grid displays objects that belong to the group icon that is currently selected in the object tree. The grid is dynamic and constantly updates its contents to reflect changes to objects.

group icon

The icon that represents a specific object group.

GUI

Graphical User Interface.

host

A machine or system.

launch

To start a task or open a window.

LUN

Logical Unit Number. Each disk in an array has a LUN. Disk partitions may also be assigned a LUN.

main window

The main VERITAS Volume Manager window. This window contains a tree and grid that display volumes, disks, and other objects on the system. The main window also has a menu bar, a toolbar, and an optional Command Launcher.

menu

A list of options or tasks. A menu item is selected by pointing to the item and clicking the mouse.

menu bar

A bar that contains a set of menus for the current window. The menu bar is typically placed across the top of a window.

mirror

A copy of a volume and its data. There can be several mirrors per volume. The terms *mirror* and *plex* are used synonymously.

node

In the VxVM tree, a node is an element attached to the tree.

In a cluster environment, a node is a host machine in a cluster.

object group

A group of objects of the same type. Each object group has a group icon and a group name. In VxVM, object groups include disk groups, disks, volumes, controllers, free disk pool disks, uninitialized disks, and file systems.

object tree

A dynamic hierarchical display of VERITAS Volume Manager objects and other objects on the system. Each node in the tree represents a group of objects of the same type.

Object View Window

A window that displays a graphical view of the volumes, disks, and other objects in a particular disk group. The objects displayed in this window are automatically updated when object properties change. This window can display detailed or basic information about volumes and disks.

Organization principle for disk group

An organization principle is associated with a disk group. These disk groups have storage pools present on them. Volumes created on such disk groups are maintained by ISP. Traditional volumes cannot be created on these disk groups. You have the option of not associating a organization principle to a disk group.

plex

A copy of a volume and its data. There can be several plexes per volume. The terms *mirror* and *plex* are used synonymously.

popup menu

A context-sensitive menu that only appears when you click on a specific object or area.

properties window

A window that displays detailed information about a selected object.

radio buttons

A set of buttons used to select optional settings. Only one radio button in the set can be selected at any given time. These buttons toggle on or off.

rules

Rules specify the storage allocation criteria. These define storage selection and storage layout. Rules are simple sentences written in the ISP language.

scroll bar

A sliding control that is used to display different portions of the contents of a window.

Search window

The VERITAS Volume Manager search tool. The Search window provides a set of search options that can be used to search for objects on the system.

splitter

A bar that separates two panes of a window (such as the object tree and the grid). A splitter can be used to adjust the sizes of the panes.

status area

An area of the main window that displays an alert icon when an object fails or experiences some other error.

storage pool - data pool, clone pool

Storage pools are associated with disk groups. Storage pools are defined by templates, capabilities and policies. Policies governing a storage pool decide

- ◆ to grow a storage pool to accommodate additional storage
- ◆ to use templates not currently associated with it.

There are two types of storage pools: data pools and clone pools. Data pools are created first. All other storage pools that are created within a disk group are clone pools.

subdisk

A set of contiguous disk blocks that form a logical disk segment. Subdisks are associated with plexes (mirrors) to form volumes.

task properties window

A window that displays detailed information about a task listed in the Task Request Monitor window.

task request monitor

A window that displays a history of tasks performed in the current VERITAS Volume Manager session. Each task is listed with the task originator, the task status, and the start/finish times for the task.

toolbar

A set of buttons used to access VERITAS Volume Manager windows. These include another main window, a task request monitor, an alert monitor, a search window, and a customize window.

tree

A dynamic hierarchical display of objects on the system. Each node in the tree represents a group of objects of the same type.

uninitialized disks

Disks that are not under VxVM control.

user templates

User templates are a collection of capabilities, templates and rules defined by the user to simplify tasks. You can group the capabilities you want the volume to have, and save these capabilities along with the rules as a user template.

volume

A virtual disk or entity that is made up of portions of one or more physical disks.

Volume Layout Details Window

A window that displays a graphical view of a volume and its components. The objects displayed in this window are not automatically updated when the volume properties change.

Volume to Disk Mapping Window

A window that displays a tabular view of volumes and their underlying disks. This window can also display details such as the subdisks and gaps on each disk.

volume set

The volume set feature provides support for the multi-device file system enhancement of VERITAS File System (VxFS). A volume set allows several volumes to be treated as a single object with one logical I/O interface.

VxFS

VERITAS File System.

VxVM

VERITAS Volume Manager.

Index

A

- activation mode 250
- add storage pool disks 237
- addby attribute 147
- adding
 - disks 42
 - file systems 167, 172
- adding a foreign device 215
- adding a mirror to a volume 106
- adding a new attribute 243
- adding a volume to a volume set 127
- adding columns 102
- alerts
 - icons 6
 - on disks 258
 - on volumes 262
- allocation unit 168
- annotating a disk 241
- annotation service
 - adding attributes 243
 - delete a short-cut 246
 - modifying attributes 244
 - remove attribute 244
 - sorting by attribute 245
- application volumes 223
- associate template set 239
- attribute short-cut 246
- attributes 241

B

- backing up file systems 174
- backups
 - creating for volumes 133
- block size 168
- boot disk
 - mirroring 47

C

- capabilities

- DCOLogStriping 280
- Raid5LogStriping 283
- capabilities options 37
- capability selection page
 - selecting a capability 226, 230
- capability tree 226, 227
- changing volume layouts 103
- checking file systems 176
- checkpoint quotas
 - disabling 189
 - enabling 188
- clearing primary policies 253
- client 12
- cluster file system
 - unmounting 182
- cluster nodes
 - cluster mounting file system on 171
- clusters
 - creating shared disk groups 249
 - importing disk groups 251
- command log 26
- commands, viewing 24
- concatenated volumes 82, 89, 136, 141
- controllers
 - disabling 58
 - enabling 58
- creating
 - disk groups 60
 - file systems 167, 172
 - instant snapshots on ISP volumes 155
 - shared disk groups 249
 - storage volumes 92
 - volume sets 126
 - volumes 82, 86, 89, 136, 141
- creating a new user template 225
- creating a storage pool 235
- critical usage 13
- customizing table headers 5

D

- data management
 - file relocation with QoS 207
- data pools 233
- DCOLogStriping capability 280
- DCOLogStriping volume template 275
- delete storage pool 236
- delete storage pool disk 238
- delete user template 231
- deleting
 - volume sets 128
- deleting a short-cut 246
- deleting volume sets 128
- deleting volumes 115
- deporting disk groups 64
- destroying disk groups 71
- dialog information panel 36
- dialog preferences 36
- dirty region logging 108
- disable path 212
- disabling
 - controllers 58
- disabling checkpoint quotas 189
- disassociate template set 240
- disconnecting
 - disks 49
- disk
 - phaseout 36
- disk arrays, renaming 59
- disk group name, changing 63
- disk groups
 - adding disks 42
 - clearing hot relocation information 69
 - creating 60, 249
 - deporting 64
 - destroying 71
 - importing 65, 251
 - joining 76
 - moving objects between 72
 - organize 78
 - preventing access 64
 - properties 78
 - recovering volumes 66
 - removing 71
 - removing disks 52
 - renaming 63
 - restoring access 65, 251
 - splitting 72, 74
 - task roadmap 41
 - upgrading 62
- disk groups versions 62
- disk usage 13
- disks
 - adding 42
 - adding to Volume Manager 42
 - alternate boot disk 47
 - annotating 241
 - boot 47
 - changing names 59
 - disconnecting 49
 - encapsulating 43
 - forced initialization 53
 - initializing 43, 53
 - maintenance tasks 7
 - mirroring 47
 - names, changing 45
 - new 42
 - online 46
 - placing under Volume Manager control 42
 - properties 22, 56, 57, 77
 - recovering 51
 - reinitializing 53
 - removing 52
 - renaming 45
 - restoring access 46
 - setup tasks 6
 - states 258
 - task roadmap 41
 - troubleshooting 53, 258
 - Volume Manager 42
 - volume mapping 13, 19
- displaying
 - commands 24
 - disk group properties 78
 - disk properties 56, 57, 77
 - dissociated objects 17
 - file system properties 177
 - object properties 12
 - Object Properties window 22
 - objects 12
 - volume properties 122
- displaying the rules page 37
- dissociated objects, displaying 17
- dissociating instant snapshots 152
- dissociating plexes 116

- E**
- enable path 212
 - enabling checkpoint quotas 188
 - enabling controllers 58
 - encapsulating disks 43
 - enclosures, renaming 59
 - extend a volume 98
- F**
- file age
 - criteria for relocation
 - multiple-component file systems 202
 - File menu 4
 - file relocation
 - configuring 194
 - criteria
 - file age 202
 - file size 202
 - criteria for
 - setup 199
 - QoS 207
 - file size
 - criteria for relocation
 - multiple-component file systems 202
 - file system size, changing 173
 - file system task roadmap 166, 193
 - file systems
 - and volume sets 172
 - and volumes 167, 171
 - backup 174
 - checking 176
 - creating 167, 172
 - maintenance tasks 8
 - monitoring capacity 177
 - repairing 176
 - resizing 173
 - snapshot copies 174
 - task roadmap 166, 193
 - unmounting 172, 173, 183
 - free space 20
- G**
- graphical view
 - object view window 13, 14
 - volume layout details 18, 22
 - grid
 - and tree view 12, 13
 - defined 4
 - GUI functions
 - associate template set 239
 - disassociate template set 240
 - resize volumes 98
 - user templates 223
- H**
- hard limits 188
 - help 5
 - Help menu 5
 - high usage 13
 - highwater mark attribute 147
 - hosts
 - accessing multiple hosts 2
 - defined 3
 - hot relocation 67
 - clearing disk group information 69
 - hot spares 67
- I**
- importing
 - disk groups 65
 - shared disk groups 251
 - initializing disks 43, 53
 - instant snapshots
 - enabling FastResync 134
 - ISP - licensing 226
 - ISP volumes
 - adding and removing columns 102
- J**
- joining disk groups 76
- L**
- layered volumes 83
 - limits
 - hard 188
 - soft 188
 - listing a foreign device 215
 - listing supported arrays 211
 - load balancing
 - subdisk move 35
 - logs
 - adding to volumes 108
 - command log 26
 - disabling 114
 - QuickLog 178, 179
 - removing 114
 - LUNs 233
- M**
- main window
 - description of 2

- status area 6
- maintenance tasks 7
- making multiple selections 24
- managing storage pools 232
- managing user templates 223
- max grow attribute 147
- menus
 - Help 5
 - Options 4
 - preferences 4
 - types of 4
 - window 4
- mirror
 - attaching 116
- mirrored 97
- mirrored volumes 83
- mirrored_data_stripped_clones storage pool
 - set 284
- mirroring
 - disks 47
 - volumes 106
- mirrors
 - adding to volumes 106
 - removing 112
- mixed layout volumes 83
- modifying attributes 244
- modifying user templates 230
- monitoring
 - alerts 6
 - performance 20
- monitoring capacity 177
- mount file system details 169
- mounting file systems 171
- moving
 - subdisks 17
- moving a subdisk 117
- multi-component file systems
 - add volume 127
- multiple component file systems
 - analyzing for effective data relocation 204
- multiple host support 2

N

- new file system details 168
- normal usage 13

O

- object properties
 - displaying 22, 56, 57, 77, 78, 122, 177

- in grid 13
- window 22
- object tree 3
- objects
 - displaying 12
 - viewing 12
- online disks 46
- Options menu 4
- organize disk groups 78

P

- performance monitoring 20
- performance tuning
 - statistics tab display 32
- placing disks under Volume Manager
 - control 42
- plexes
 - dissociating 116
- preferences 36
 - Options menu 4
 - setting 4
- primary node
 - setting 253
- primary policy
 - clearing 253
 - setting 252
- projection 16
- properties 13
 - 122
 - disk 56, 57, 77
 - disk groups 78
 - Disk Properties window 56, 57, 77
 - File System Properties window 177
 - in grid 13
 - viewing objects
 - properties 22, 78
 - volumes 122, 177

Q

- Quality of Storage Service feature 193
- QuickLog 178
 - device 179

R

- RAID-5 96
- RAID-5 logging 108
- RAID-5 volumes 82
- Raid5LogStriping capability 283
- Raid5LogStriping volume template 276
- reattaching instant snapshots 152

- recovering
 - disks 51
 - volume sets 128
 - volumes 8, 51, 66, 110
 - recovery tasks 8
 - refreshing instant snapshots 151
 - relayout 102, 103
 - relocation of files
 - QoS feature 207
 - relocation policies
 - configuring 194
 - remote administration 1
 - removing
 - disk groups 71
 - disks 52
 - logs 114
 - volume sets 128
 - removing a foreign device 218
 - removing an attribute 244
 - removing columns layouts 102
 - removing volume sets 128
 - renaming
 - disk arrays 59
 - disk groups 63
 - disks 45
 - enclosures 59
 - volume sets 128
 - volumes 99
 - report 26
 - resize volumes 98
 - resizing
 - file systems 173
 - resizing a volume 98
 - restarting and recovering volumes 110
 - restoring volumes 152
 - roadmap, disk task 41
 - root disk
 - mirroring 47
 - rules
 - displaying the rules page 37
 - rules page
 - selecting attributes for rules 227
- S**
- scanning disks 213
 - search 26
 - search feature 26
 - editing and removing 31
 - setting up a query 26
 - viewing previous search results 31
 - security 12
 - selecting a capability 226, 230
 - selecting attributes 227
 - selecting objects 24
 - selections
 - multiple 24
 - server 12
 - setting primary node 253
 - setting primary policies 252
 - setting up your system 12
 - setting usage threshold 177
 - setup tasks 6
 - shared disk groups 249, 251
 - snapshot
 - synchronization 154
 - snapshot copies of file systems 174
 - snapshot copies of volumes 133, 156
 - snapshots
 - copy of file system 174
 - copy of volume 133, 156
 - dissociating 152, 160
 - enabling FastResync 134
 - merging 159
 - reattaching 152
 - refreshing 151
 - restoring volumes 152
 - splitting 154
 - soft limits 188
 - sorting by attribute name 245
 - splitter 4
 - splitting
 - subdisks 119
 - splitting disk groups 72, 74
 - splitting instant snapshot hierarchy 154
 - splitting subdisks 119
 - start synchronization 154
 - starting VEA 12
 - statistics
 - viewing 34
 - statistics tab view 32
 - status area 6
 - stopping volumes 110
 - storage pool policy 233
 - storage pool sets
 - mirrored_data_stripped_clones 284
 - storage pools
 - add disks 237
 - creating 235

- delete 236
- delete disks 238
- managing 232
- remove disks 238
- storage volumes
 - creating 92
- striped 95
- striped volumes 82
- subdisk move command 117
- subdisks
 - graphical view 18, 22
 - moving 17, 117
 - numbering 119
 - splitting 119
 - viewing 18, 22
- synchronization start 154
- system
 - setup 12

T

- table headers
 - customizing 5
- task dialogs 36
- tasks
 - disk maintenance 7
 - disk setup 6
 - displaying 24
 - file system maintenance 8
 - list 6
 - maintenance 7
 - properties 26
 - recovery 8
 - repair 8
 - request monitor window 24
 - roadmap 6
 - setup 6
 - throttling 25
 - viewing 24
 - volume maintenance 8
 - volume recovery 8
 - volume repair 8
 - volume setup 6, 7
- tear off 5
- templates
 - DCOLogStriping 275
 - Raid5LogStriping 276
- threshold percentage 177
- toolbar 5
- tree

- and grid view 12, 13
- object 3
- troubleshooting
 - disk states 258
 - volume states 262

U

- unmounting file systems 172, 173, 183
- usage legend 13
- user templates
 - creating new user templates 225
 - delete 231
 - modifying 230
- using the search feature 26

V

- VEA
 - starting 12
- view preferences 36
- viewing commands 24
- viewing objects
 - and object properties 12
 - graphical view 18, 22
 - properties 56, 57, 77, 122, 177
 - tree and grid 13
- viewing real-time statistics 34
- viewing tasks 24
- Volume Layout Details window
 - described 13, 18, 22
- volume layouts 82, 103
- volume name, changing 99
- volume set
 - recovering 128
- volume set name, changing 128
- volume sets
 - adding a volume 127
 - adding file systems 172
 - and file systems 172
 - creating 126
 - deleting
 - removing 128
 - removing a volume 129
 - renaming 128
 - volume states 262
- volume task roadmap 80, 125, 132
- volume templates
 - DCOLogStriping 275
 - Raid5LogStriping 276
- volume to disk mapping window 13, 19
- volumes 86

- adding file systems 167
- adding logs 108
- adding mirrors 106
- and file systems 167, 171
- attaching mirrors 116
- backing up 133
- backup and restore 111, 163
- changing layout 103
- concatenated 82, 89, 136, 141
- concatenated pro 83
- creating 82, 86, 89, 136, 141
- deleting volumes 115
- disabling 110
- disk mapping 13, 19
- dissociating plexes 116
- dissociating snapshots 160
- enabling FastResync 134
- extending 98
- graphical view 13, 18, 22
- layered 83
- layout 13
- maintenance tasks 8
- merging snapshots 159
- mirrored 83, 97
- mixed layout 83
- mounting file systems 171
- preparing to restore 111, 163
- properties 22, 122, 177
- RAID-5 82, 96
- recovering 51, 66, 110
- recovering/restarting 110
- recovery tasks 8
- relayout 102, 103
- removing from a volume set 129
- removing logs 114
- removing mirrors 112
- renaming 99
- repair tasks 8
- resizing 98
- setting usage 103
- setup tasks 6, 7
- snapshot copies 133, 156
- splitting subdisks 119
- stopping 110
- striped 82, 95
- striped pro 83
- task roadmap 80, 125, 132
- troubleshooting 262
- underlying disks 13, 19
- unmounting file systems 172, 173, 183
- VRTSexplorer xxii
- VxVM
 - obtaining system information xxii

W

- window menu 4
- window preferences 36
- windows
 - graphical view 18, 22
 - main window 2
 - object properties 13, 22
 - object view 14
 - resizing 4
 - volume layout details 18, 22
 - volume to disk mapping window 19