

SCO Forum 2006

MOBILITY EVERYWHERE >



Presentation Title: Protect and Manage SCO OS & Data using ODM & Mirroring
Presenter's Name: Dave Stetzel
Session ID: 111 & 127

1



Platinum Sponsor



Get Your Passport Stamped



- Be sure to get your Passport stamped.
 - Get your passport stamped
 - By breakout session instructors
 - By exhibitors in the exhibit hall
 - Turn in your Passport
 - After the last breakout session on Wednesday
 - Drawing for great prizes for Wrap-up Session
- Remember to complete the breakout session evaluation form, too

WIN BIG

SCO Forum 2006
PASSPORT

Turn in this card at the Registration and Information desk. Prize drawings will be held during the Closing Session of SCO Forum, at 4pm on Tuesday, August 9th. You must be present to win.

HOW
> Att
> Visi
> Hav
Atten
a drap
iPods

Name: _____
Company: _____
eMail: _____
Phone: _____


Breakout Sessions

Monday: ○ ○ ○
Tuesday: ○ ○ ○

Tradeshow

○ ○ ○ ○ ○ ○ ○ ○ ○ ○

SCO Automates Transactions



SCO UNIX Products



Agenda



- What is ODM and Mirroring
- Model Numbers and Pricing
- Documentation Resources
- ODM Technology Overview
- RAID Levels Supported
- ODM Product Capabilities
- Initializing the Volume Manager
- HowTo Mirror boot (root) drive
- Recovery
- Summary
- Q&A



What is ODM & Mirroring?

What is ODM & Mirroring?



Low cost Data Manager which provides customers with the tools to easily manage and protect their business critical data on OpenServer 6 and UnixWare 7 systems.

- Both the Online Data Manager (ODM) and Mirroring can be coupled with any standard off the shelf disk system, including the new low cost Serial ATA (SATA) drives.
- Using these new tools with low cost disk systems will significantly lower the cost of solutions requiring easy to use data management and highly reliable data security.

What is ODM & Mirroring?



Products

- Online Data Manager (ODM)
 - SCO Online Data Manager provides software RAID Levels 0, 1, 5, 10
 - It includes a graphical Visual Administrator for exceptional ease-of-use.
 - Supports online filesystem resizing, Online disk performance analysis, migration, and RAID level modifications. ODM is a cost-effective, enterprise-class storage management solution for high availability and online volume management.

What is ODM & Mirroring?



Products continued:

- **Disk Mirroring**
 - Customers who just need software disk mirroring solution to enhance the availability of disk storage systems without any other RAID features, have the option of purchasing the Mirroring license.
 - The Mirroring option provides increased data availability by providing fault tolerance against disk failures and faster data access.
 - By adding a second disk to the system and adding Mirroring, administrators will automatically have access to their business-critical data should one disk fail.



Model Numbers and Pricing



OpenServer 6

Data Management	License Pack	Web License	US\$ List	Euro List
OpenServer 6 Online Data Manager License	LA215-UX00-6.0	LA215-UX00W-6.0	\$349	€ 303
OpenServer 6 Disk Mirroring License	LA216-UX00-6.0	LA216-UX00W-6.0	\$99	€ 86

The required software is included on the OpenServer 6 Maintenance Pack 2 CD and available on our web download page.



UnixWare 7.1.4

Data Management	License Pack	Web License	US\$ List	Euro List
Online Data Manager 3.2 License	LA415-UW70-3.2	LA415-UW70W-3.2	\$349	€ 303
Disk Mirroring 3.2 License	LA416-UW70-3.2	LA416-UW70W-3.2	\$99	€ 86

The required software is included on the Base Operating System CD and on the SCO Support download web page as part of the *recut* CD1 media.



Documentation Resources



Data Sheet

- <http://www.sco.com/products/optionalservices/odm.html>

Product Documentation *(Must Read Documents!)*

- Online Data Manager overview and installation
 - Provides information about licensing ODM products, software notes and recommendations, and procedures for initializing the VERITAS Volume Manager.



Product Documentation, continued

- VERITAS Volume Manager User's Guide
 - Description and Overview of Volume Manager
 - Provides basic procedures for volume management and system administration
 - Use of the command line (CLI) utilities
 - Use of the Menu Interface Operations facility
 - Use of the Volume Manager Visual Administrator (VxVA), the graphical interface to ODM



Product Documentation, continued

- VERITAS Volume Manager System Administrator's Guide
 - Provides information about VxVM configuration
 - Performance monitoring
 - Disk and Disk Group administration
 - Volume administration
 - Recovery procedures
 - Includes VxVM error messages
- Manuals are loaded as part of the ODM product installation and can be viewed and/or printed via the DocView facility.



ODM Technology Overview



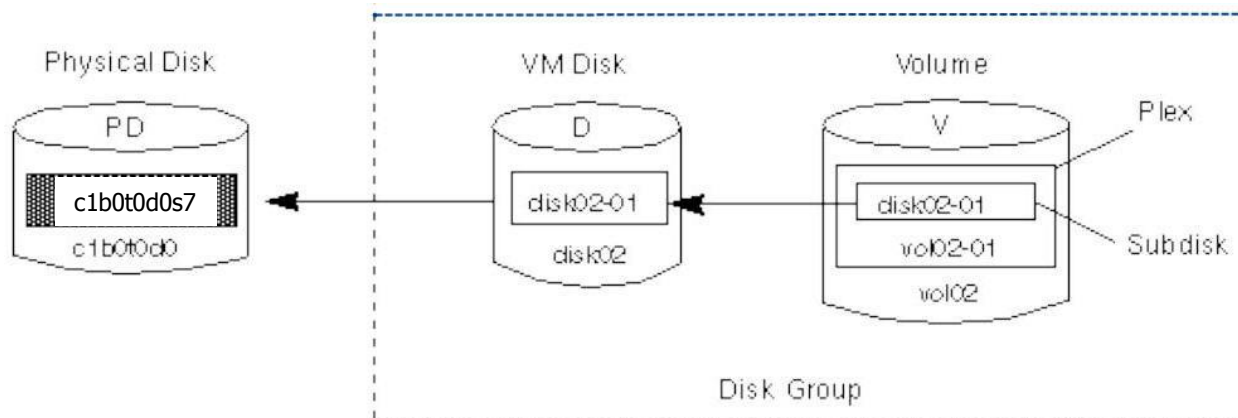
Relationship of VxVM Objects

Physical Objects

Volume Manager Objects

Physical Disk (PD)

- Partitions (Slices)



Volume Manager

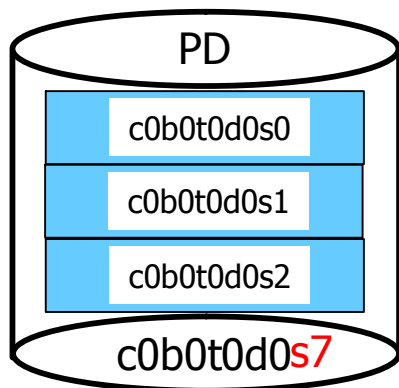
- VM disks
- Disk groups
- Subdisks
- Plexes
- Volume



Object Definitions - Physical

Physical Disk (PD) - physical storage device (media), which can be accessed using a device name *c#b#t#d#*

Partitions - used synonymously with division or slice. The *s#* at the end of the device name, *c#b#t#d#s#*



Example - OpenServer 6 Boot disk

-s0 stand slice

-s1 swap slice

-s2 root slice

-s7 entire disk

Volume Manager Object Definition




Object Definitions – Volume Manager

VM Disk (D) -

The contiguous area of physical disk space assigned to Volume Manager to allocate space from. The VM disk is accessed using a unique *disk media name*.

The VM disk is composed of a **public region** (space allocation area), and **private region** (where configuration information is stored).



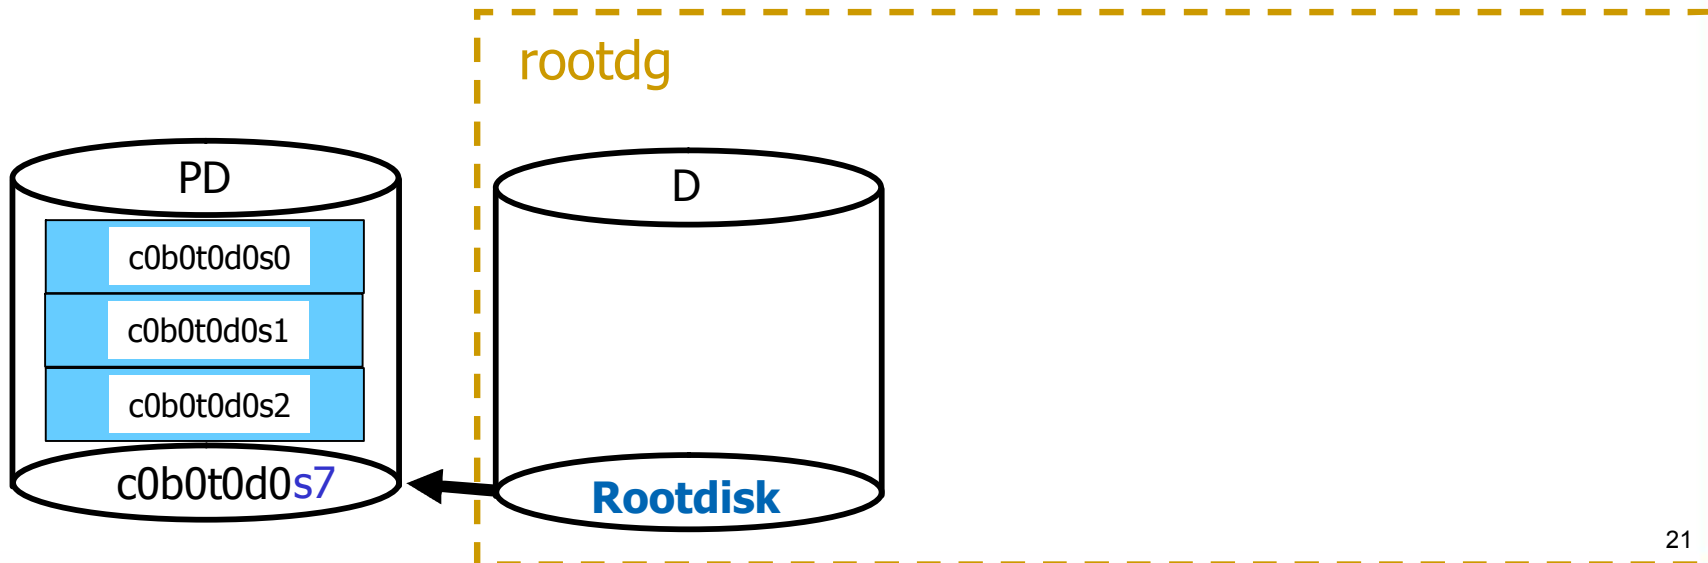
Volume Manager Object Definition



Object Definitions – Volume Manager

Disk Group -

Collection of VM disks that share a common configuration. The default disk group is **rootdg** (the root disk group). Allow the administrator to group disks into logical sets. Volumes are created in disk groups.



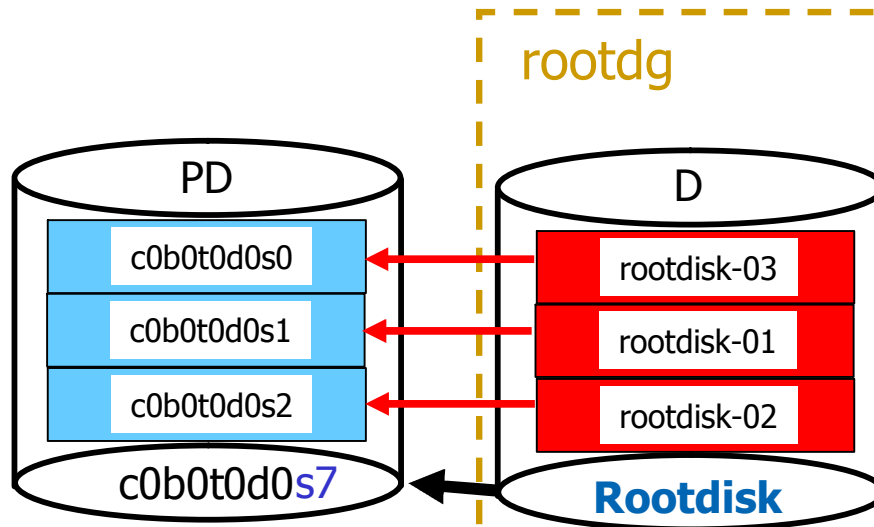
Volume Manager Object Definition



Object Definitions – Volume Manager

Subdisks -

A VM Disk can be divided in one or more **subdisks**. Each subdisk represents a specific portion of a VM disk, which is mapped to a specific region of a physical disk.



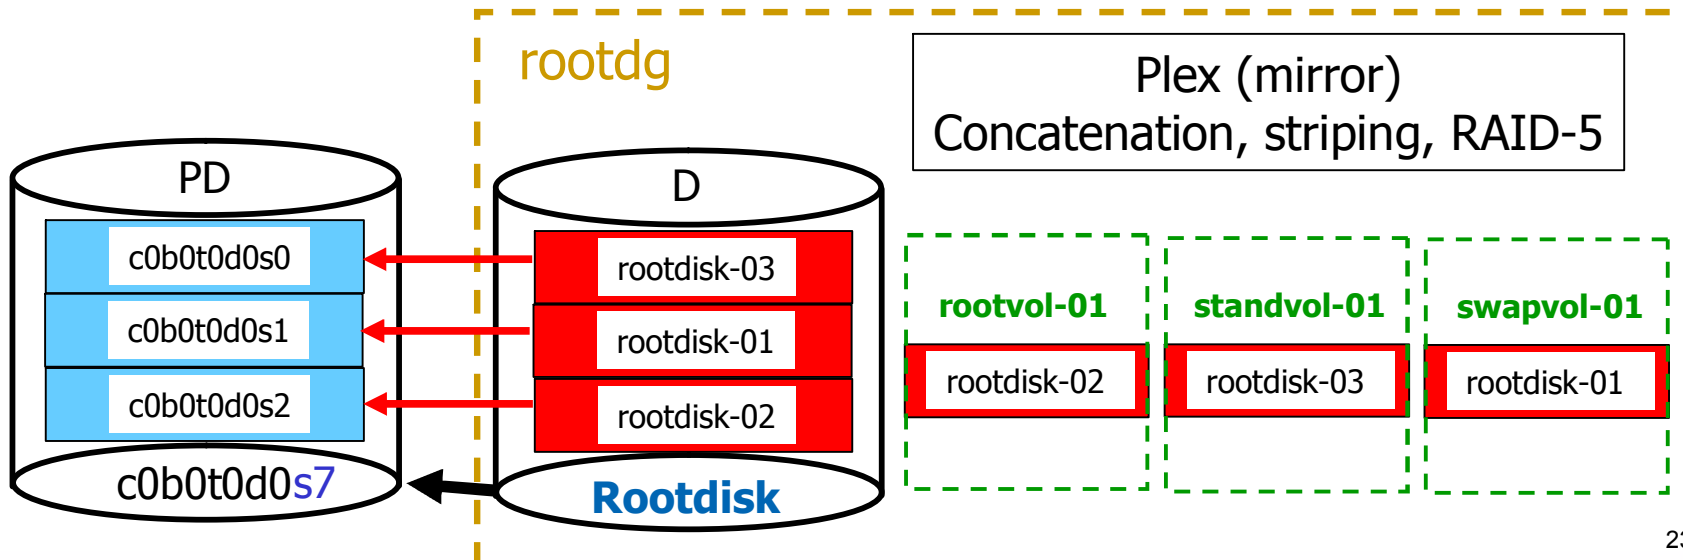
Volume Manager Object Definition



Object Definitions – Volume Manager

Plexes -

A **plex** is a virtual entity made up of one or more subdisks on one or more disks and in default is a mirror. In addition to its default mirror, a plex can be organized on subdisks in three ways:

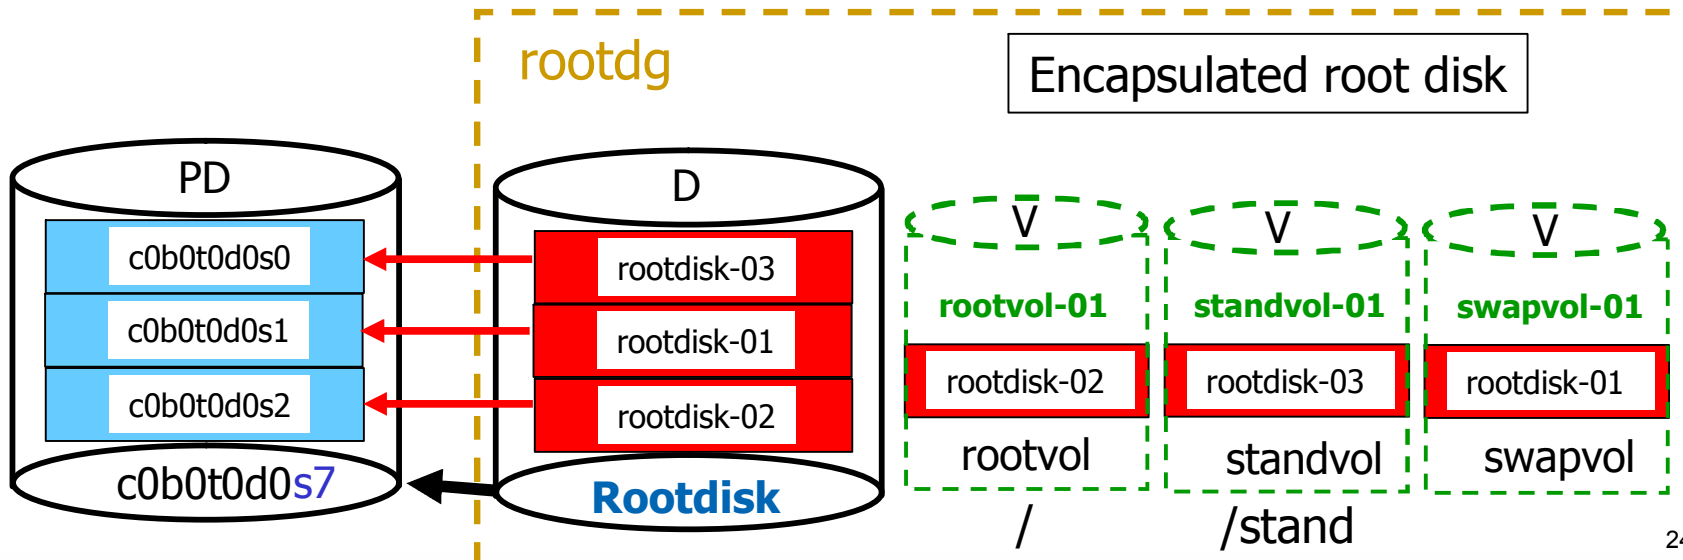


Volume Manager Object Definition



Object Definitions – Volume Manager

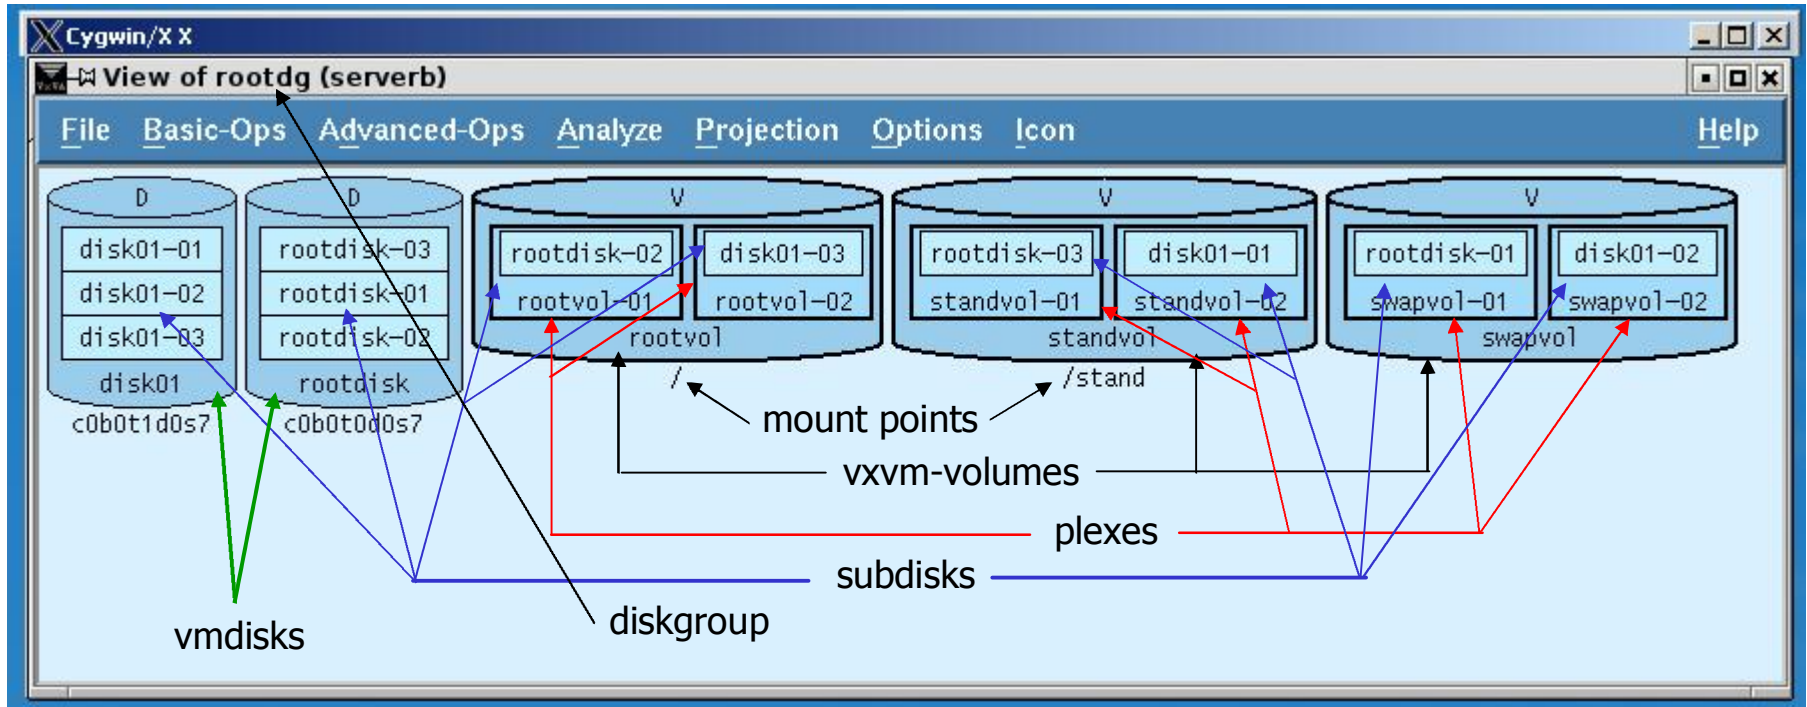
Volume - A volume is a virtual disk device that appears like a physical disk division, but does not have the physical limitations of a physical disk division. The configuration of a volume can be changed without causing disruption.



Object Relationship



View of rootdg for mirrored root disk





Device Nodes, *example*: root slice (OSR6)

- Physical Disks
 - /dev/dsk/c0b0t0d0s2
 - /dev/rdisk/c0b0t0d0s2
- VxVM Volumes
 - /dev/vx/dsk/rootvol
 - /dev/vx/rdisk/rootvol
- /etc/vfstab
 - Encapsulation of existing partitions will cause file system entries in vfstab file to be modified to VxVM Volumes. Except for encapsulated root volume.



Volume Manager Supports the following user interfaces:

- Visual Administrator Interface (SCOadmin)
 - Graphical user interface
 - Provides visual elements such as icons, menus, and forms to ease the task of manipulating Volume Manager objects
- Command Line Interface (CLI)
 - Comprehensive set of commands that range from simple minimal user input commands to complex, requiring detailed user input.
- Support (Menu) Operations Interface (vxdiskadm)
 - Provides menu-driven interface for performing disk and volume manager functions.



RAID Levels Supported



Volume Manager supports the following RAID levels:

- RAID-0 (Striping)
 - Mapping of data so that the data (called stripe units) is interleaved among two or more physical disks. The default stripe unit size is 64 kilobytes.
 - The striped plex subdisks are grouped into columns, with each physical disk limited to one column. Each column can contain one or more subdisks from one or more physical disks.
 - Can provide increased performance, but single disk failure will cause the volume to fail.

Reference pg 12 VxVM Administrator's Guide



Volume Manager supports the following RAID levels:

- RAID-1 (Mirroring)
 - Duplicating of data contained in a volume using multiple mirror plexes.
 - If a physical disk failure occurs, the mirror on the failed disk becomes unavailable, but the system continues to operate using the unaffected mirrors.



Volume Manager supports the following RAID levels, cont.:

- RAID-0 plus RAID-1 (Striping and Mirroring)
 - Combination of striping with mirroring.
 - When used together on the same volume, offers the benefits of spreading data across multiple disks while providing redundancy protection.



Volume Manager supports the following RAID levels, cont.:

- RAID-5 Array
 - Provides data redundancy through the use of parity.
 - As data is being written to a RAID-5 Volume, parity is also calculated and the resulting parity is written to the volume.
 - If a portion of a RAID-5 volume fails, the data that was on that portion of the failed volume can be reconstructed from the remaining data and parity.
 - Refer to documentation for various RAID-5 implementation schemes.



Logging

- It is possible for data not committed to an active write to be lost or corrupted if a disk and system fail together. Therefore, if this double-failure occurs, the recovery of the corrupted disk may be corrupted itself.
- Logging is used to prevent this corruption of recovery data. A log of the new data and parity is kept on a persistent device. The new data and parity are then written to the disks.
- Logs are associated with RAID-5 volume by being attached as non-RAID-5 layout plexes. More than one log plex can exist per RAID-5 volume, which mirrors the log areas.



Hot-Relocation automatically reacts to I/O failures on redundant (mirrored or RAID-5) objects to restore redundancy.

- When Volume Manager hot-relocation daemon detects I/O failures on VxVM objects it relocates the effected subdisks to disks designated as spare disks and/or available space within the disk group.
- A successful hot-relocation process involves:
 1. Detecting events resulting from failure of a disk, plex, RAID-5, or mirror subdisk.
 2. Notifying the system administrator (and designated users) of the failure and identifying the affected VxVM objects, via electronic mail.
 3. Determines which subdisks can be relocated, finds space for the subdisks, and relocating the subdisks. Notifies system administrator of actions.
 4. Initiates any necessary recovery procedure to restore data. Notifies the system administrator of the recover outcome.



ODM Product Capabilities

Menu (Disk) Operations



vxdiskadm

```
192.168.0.9 - PuTTY
Volume Manager Support Operations
Menu: VolumeManager/Disk

1      Add or initialize one or more disks
2      Encapsulate one or more disks
3      Remove a disk
4      Remove a disk for replacement
5      Replace a failed or removed disk
6      Mirror volumes on a disk
7      Move volumes from a disk
8      Enable access to (import) a disk group
9      Remove access to (deport) a disk group
10     Enable (online) a disk device
11     Disable (offline) a disk device
12     Mark a disk as a spare for a disk group
13     Turn off the spare flag on a disk
list   List disk information

?      Display help about menu
??     Display help about the menuing system
q      Exit from menus

Select an operation to perform: █
```

File System Operations



Visual Administrator performs file system operations by executing file system commands directly or in combination of vx commands.

- Create file system on volume (*mkfs*)
 - Simple volume
 - Striped Volume
 - RAID 5 Volume
 - Mirrored Volume
- Mount and unmount (*mount, umount*)
- Resize, Defragment, (*fsadm*) Snapshot (*vxassist, vxedit*)
 - grow and shrink file systems (and volumes) (*fsadm, vxassist*)
 - create snapshot mirror volume (*vxassist, vxedit*)
 - Mirror volume, mounts read only for online backup of volumes
 - To maintain file system performance, *fsadm* should be run periodically against all VxFS file systems to reduce fragmentation.



vxrelacd(ADM)

- vxrelacd - monitors the Volume Manager for failure events and relocates failed subdisks
- Mail notification
 - By default, vxrelacd sends electronic mail to root with information about a detected failure and the status of any relocation and recovery attempts.
 - To instruct vxrelacd to send this mail (via mailx) to other users, add the desired user logins to the vxrelacd startup line in the startup script `/etc/rc2.d/s95vxvm-recover` and then reboot the system.
- `/etc/init.d/vxvm-recover`
 - `/usr/lib/vxvm/bin/vxrelacd root scouser &`



VxVM Performance Monitoring

- VxVm Provides two types of performance information
 - I/O statistics
 - Retrieved using the vxstat utility.
 - Provides information for activity on volumes, plexes, subdisks, and disks under VxVM control.
 - I/O traces
 - The vxtrace utility is used to trace operations on volumes.
 - Prints kernel I/O error or I/O trace records.
- Performance Data gathered can then be used to determine if reconfiguration is needed for optimum system performance.



Low-End Server Benchmark Tests, provided by iXorg

- LEAD:
- =====
- Steve James

- DESCRIPTION:
- =====
- Evaluate four low-end servers for performance comparison with no disk redundancy, H/W disk mirror redundancy, and S/W disk mirror redundancy.

- PROCEDURES:
- =====
- H/W Mirroring using Adaptec 1210SA SATA RAID-1 (from SCO CHWP and TA #126029).
- S/W Mirroring using SCO ODM Mirror.



Low-End Server Benchmark Tests, provided by iXorg

- Run following commands and record results:
- Backup the entire OSR6 installed configuration...
- #1 edge cvbf 64 /dev/null .
- #2 edge -zBUFFERS=80 -cSbf 256 /dev/null .
- #3 edge cbf 64 /dev/null .
- Create one 1GB file in /test directory, then cd /test...
- #4 edge cvbf 64 /dev/null .
- #5 edge -zBUFFERS=80 -cSbf 256 /dev/null .

Performance Monitoring cont.



Low-End Server Benchmark Tests, provided by iXorg

- TECHNICAL ISSUES:

- =====

- The Adaptec SATA RAID controller did not work in either JBOD or RAID-1 mode under OpenServer 6.0.0.

- RESULTS:

- =====

- Machine Configurations:

Pseudonym	CPU	RAM	Disk
=====	=====	====	
=====			
AMD64	AMD Athlon64 3.0GHz	1GB	100MB/sec IDE
AMD64B	AMD Sempron ?GHz	1GB	300MB/sec SATA
MG's Intel	Intel P4 w/ HT 3.0GHz	1GB	150MB/sec SATA
HP ML310	Intel Dual-Core 3.4GHz	1GB	150MB/sec SATA

Performance Monitoring cont.



Low-End Server Benchmark Tests, provided by iXorg

Machine	No ODM	With ODM Mirror
=====	=====	
=====	=====	
AMD64		
#1	03:07 = 53.174 GB/hour	Not Run
#2	03:08 = 52.784 GB/hour	
#3	03:05 = 53.726 GB/hour	
#4	00:18 = 186.555 GB/hour	
#5	00:17 = 189.206 GB/hour	



Low-End Server Benchmark Tests, provided by iXorg

- AMD64B

- #1 03:19 = 49.709 GB/hour 03:41 = 82.563 GB/hour
- #2 02:54 = 56.618 GB/hour 03:38 = 83.700 GB/hour
- #3 02:54 = 56.680 GB/hour 03:37 = 84.060 GB/hour
- #4 00:17 = 205.184 GB/hour 00:16 = 206.940 GB/hour
- #5 00:16 = 211.464 GB/hour 00:16 = 208.341 GB/hour

- MG's Intel

- #1 02:32 = 38.334 GB/hour 03:23 = 67.278 GB/hour
- #2 02:52 = 33.877 GB/hour 03:24 = 67.125 GB/hour
- #3 02:54 = 33.460 GB/hour 03:20 = 68.286 GB/hour
- #4 00:17 = 207.437 GB/hour 00:16 = 220.300 GB/hour
- #5 00:17 = 214.149 GB/hour 00:16 = 220.732 GB/hour



Low-End Server Benchmark Tests, provided by iXorg

- HP ML310
- #1 03:59 = 42.770 GB/hour 03:24 = 49.403 GB/hour
- #2 03:24 = 50.080 GB/hour 03:16 = 51.261 GB/hour
- #3 03:24 = 50.082 GB/hour 03:15 = 51.550 GB/hour
- #4 00:18 = 194.459 GB/hour 00:18 = 196.491 GB/hour
- #5 00:18 = 191.057 GB/hour 00:18 = 195.809 GB/hour

- CONCLUSIONS:
- =====
- 1. Using SCO S/W Mirror boosted disk subsystem performance by up to 104% for backups of many small files.
- 2. Using SCO S/W Mirror boosted disk subsystem performance for backups of very large files, but not as much as for many small files.



Initializing the Volume Manager

Initializing the Volume Manager



This section covers steps necessary to initialize the Volume Manager. Reference Overview & Installation manual.

1. Install latest MP & Verify (If not already installed)
 - MP2 for OpenServer 6 & MP3 for UnixWare 7.1.4
 - <http://sco.com/support/update/download>

Initializing the Volume Manager



2. Plan configuration

- Know the system configuration, controller(s), channels, & drives
- Drives to be put under Volume Manager control
- Drives to be encapsulated
- Drives to be mirrored (Master and Target(s))
 - Target mirror drive size is \geq master drive
 - Plan for Boot after failure (VxVM Boot Floppy)
- RAID configuration, 1, 0&1, 5
- Adjust configuration for VTOC changes if needed
- Condition drives if needed (mkdev hd)

Initializing drives for target mirrors recommend creating a active 100% UNIX partition and create one slice (division) on the partition, - write the partition table. Also, if drive will be used as possible boot mirror, "Overwrite system master boot code."

Initializing the Volume Manager



Plan, configuration being used for this training:

```
192.168.0.9 - PuTTY
bash-3.1# sdiconfig -l
0:0,7,0: HBA      : (adsb,1) Adaptec PCI SCSI
  0,0,0: DISK     : IBM      DDRS-34560D      DC1B
  0,1,0: DISK     : IBM      DDRS-34560D      DC1B
  0,2,0: DISK     : IBM      DCHS04W          2727
  0,3,0: DISK     : IBM      DCHS04W          6464
  0,4,0: CDROM    : NEC      CD-ROM DRIVE:4651.03
  0,5,0: CDROM    : YAMAHA   CRW4416S         1.0g
  0,6,0: TAPE     : ARCHIVE  Python 25501-XXX2.34
1:0,6,0: HBA     : (c8xx,1)Symbios 875-3 42400
bash-3.1# █
```

Initializing the Volume Manager



Default vtoc of boot drive:

slice 0: STAND	permissions: VALID	starting sector: 32067 (cyl 1)	length: 81918 (5.10 cyls)
slice 1: SWAP	permissions: VALID UNMOUNTABLE	starting sector: 113985 (cyl 7)	length: 880638 (54.82 cyls)
slice 2: ROOT	permissions: VALID	starting sector: 994623 (cyl 61)	length: 7921446 (493.09 cyls)
slice 7: DISK	permissions: VALID UNMOUNTABLE	starting sector: 16065 (cyl 1)	length: 8900010 (554.00 cyls)
slice 13: BOOT	permissions: VALID UNMOUNTABLE	starting sector: 16065 (cyl 1)	length: 16002 (1.00 cyls)

VxVM requires two special vtoc slices:

Slice 14 (0x0e)

- Volume Management Public Slice
(Area of disk under VxVM control)

Slice 15 (0x0f)

- Volume Management Private Slice
(Configuration data storage, usually 1 cyl in size but depends on system VxVM configuration)

Initializing the Volume Manager



3. Install ODM, license and patch

- Mirror Only License or Full ODM License
- Oss706a (OpenServer 6)

4. Verify PATH environment variable

- See page 7 ODM Overview & Installation

```
/usr/sbin:/etc/vx/bin:/opt/vxvm-va/bin
```

5. If you have any disks that you do not want under Volume Manager control create the file */etc/vx/disks.exclude* containing the disks vxinstall is to ignore:

- *c0b0t2d0*
- *c0b0t3d0*

Initializing the Volume Manager



6. As root user initialize the Volume Manager

- ODM Overview & Installation
- Start Initialization

vxinstall

```
192.168.0.9 - PuTTY
Volume Manager Installation
Menu: VolumeManager/Install

The Volume Manager names disks on your system using the controller
and disk number of the disk, substituting them into the following
pattern:

    c<controller>b<bus>t<target>d<disk>

Some examples would be:

    c0b0t0d0      - first controller, bus 0, first target, first disk
    c1b0t0d0      - second controller, bus 0, first target, first disk
    c1b0t1d0      - second controller, bus 0, second target, first disk

The Volume Manager has detected the following controllers on your system:

    c0:

Hit RETURN to continue. █
```

Initializing the Volume Manager



```
192.168.0.9 - PuTTY
Volume Manager Installation
Menu: VolumeManager/Install

You will now be asked if you wish to use Quick Installation or
Custom Installation. Custom Installation allows you to select how
the Volume Manager will handle the installation of each disk
attached to your system.

Quick Installation examines each disk attached to your system and
attempts to create volumes to cover all disk partitions that might
be used for file systems or for other similar purposes.

If you do not wish to use some disks with the Volume Manager, or if
you wish to reinitialize some disks, use the Custom Installation
option. Otherwise, we suggest that you use the Quick Installation
option.
Hit RETURN to continue. █
```

Initializing the Volume Manager



Prompt for quick or custom install:

- Quick is good to encapsulate root drive in group "rootdg" and default initialization of disks.
- Custom enables each disk to be manipulated individually.

```
192.168.0.9 - PuTTY
Volume Manager Installation Options
Menu: VolumeManager/Install

1      Quick Installation
2      Custom Installation

?      Display help about menu
??     Display help about the menuing system
q      Exit from menus

Select an operation to perform: 1
```

Initializing the Volume Manager



VxVM install recognizes Boot Disk and gives option to encapsulate:

```
192.168.0.9 - PuTTY
Volume Manager Custom Installation
Menu: VolumeManager/Install/Custom

The c0b0t0d0 disk is your Boot Disk.  You can not add it as a new
disk.  If you encapsulate it, you will make your root filesystem
and other system areas on the Boot Disk into volumes.  This is
required if you wish to mirror your root filesystem or system
swap area.

Encapsulate Boot Disk [y,n,q,?] (default: n) y
```

Initializing the Volume Manager



Prompts for disk name of Boot Disk, suggest use the default name *rootdisk*:

```
192.168.0.9 - PuTTY
Volume Manager Custom Installation
Menu: VolumeManager/Install/Custom

The c0b0t0d0 disk is your Boot Disk. You can not add it as a new
disk. If you encapsulate it, you will make your root filesystem
and other system areas on the Boot Disk into volumes. This is
required if you wish to mirror your root filesystem or system
swap area.

Encapsulate Boot Disk [y,n,q,?] (default: n) y
Enter disk name for c0b0t0d0 [<name>,q,?] (default: rootdisk) █
```


Initializing the Volume Manager



Verifies that you want to continue with encapsulation of Boot Disk:

```
192.168.0.9 - PuTTY
Volume Manager Custom Installation
Menu: VolumeManager/Install/Custom

The c0b0t0d0 disk is your Boot Disk.  You can not add it as a new
disk.  If you encapsulate it, you will make your root filesystem
and other system areas on the Boot Disk into volumes.  This is
required if you wish to mirror your root filesystem or system
swap area.

Encapsulate Boot Disk [y,n,q,?] (default: n) y

Enter disk name for c0b0t0d0 [<name>,q,?] (default: rootdisk)

The c0b0t0d0 disk has been configured for encapsulation.

Hit RETURN to continue. █
```

Initializing the Volume Manager



Detection of remaining disks on controller c0:

```
192.168.0.9 - PuTTY
Volume Manager Quick Installation
Menu: VolumeManager/Install/QuickInstall/c0
Generating list of attached disks on c0....

<excluding root disk c0b0t0d0>
<excluding c0b0t2d0>
<excluding c0b0t3d0>

  The Volume Manager has detected the following disks on controller c0:

  c0b0t1d0

Hit RETURN to continue. █
```

Initializing the Volume Manager



Installation options of disks on controller c0:

```
192.168.0.9 - PuTTY
Volume Manager Quick Installation For Controller c0
Menu: VolumeManager/Install/QuickInstall/c0

Initialize all disks on this controller ? (destroys data on these disks)
[y,n,q,?] (default: n) █
```

Initializing the Volume Manager



Installation options for disk c0b0t1d0:

```
192.168.0.9 - PuTTY
Volume Manager Quick Installation For Controller c0
Menu: VolumeManager/Install/QuickInstall/c0

Initialize all disks on this controller ? (destroys data on these disks)
[y,n,q,?] (default: n)

    Volume Manager will now try to encapsulate all the disks on this controller.
    Disks not having valid partitons will be initialized.
Hit RETURN to continue. █
```

Initializing the Volume Manager



Prompts for VxVM disk naming to be used:

```
192.168.0.9 - PuTTY
Volume Manager Quick Installation
Menu: VolumeManager/Install/QuickInstall/c0/Encap
Use default disk names for these disks? [y,n,q,?] (default: y) █
```

Initializing the Volume Manager



Displays disk name c0b0t1d0:

```
192.168.0.9 - PuTTY
Volume Manager Quick Installation
Menu: VolumeManager/Install/QuickInstall/c0/Encap
Use default disk names for these disks? [y,n,q,?] (default: y)
    The c0b0t1d0 disk will be given disk name disk01
    The c0b0t1d0 disk has been configured for encapsulation.
Hit RETURN to continue. █
```

Initializing the Volume Manager



Verification screen before completing VxVM initialization:

```
192.168.0.9 - PuTTY
Volume Manager Quick Installation
Menu: VolumeManager/Install/QuickInstall

The following is a summary of your choices.

      c0b0t0d0      Encapsulate
      c0b0t1d0      Encapsulate

Is this correct [y,n,q,?] (default: y) █
```

Initializing the Volume Manager



Shutdown and reboot is required to finish Volume Manager Initialization:

```
192.168.0.9 - PuTTY
Volume Manager Quick Installation
Menu: VolumeManager/Install/QuickInstall

The following is a summary of your choices.

      c0b0t0d0      Encapsulate
      c0b0t1d0      Encapsulate

Is this correct [y,n,q,?] (default: y)

The system now must be shut down and rebooted in order to continue
the reconfiguration.

Shutdown and reboot now [y,n,q,?] (default: n) y
```


Initializing the Volume Manager



Shutdown :

```
192.168.0.9 - PuTTY
Volume Manager Quick Installation
Menu: VolumeManager/Install/QuickInstall

The following is a summary of your choices.

      c0b0t0d0      Encapsulate
      c0b0t1d0      Encapsulate

Is this correct [y,n,q,?] (default: y)

The system now must be shut down and rebooted in order to continue
the reconfiguration.

Shutdown and reboot now [y,n,q,?] (default: n) y

Shutdown started.      Sat Jul 22 17:35:34 CDT 2006

Broadcast Message from root (ttypl) on serverb   Jul 22 17:35 2006...

THE SYSTEM IS BEING SHUT DOWN NOW ! ! !
Log off now or risk your files being damaged.

Shutdown proceeding. Please wait .....
```

Initializing the Volume Manager



List disks and status following VxVM initialization:

```
192.168.0.9 - PuTTY
bash-3.1# vxdisk list
DEVICE      TYPE      DISK      GROUP     STATUS
c0b0t0d0s7  sliced   rootdisk  rootdg    online
c0b0t1d0s7  sliced   disk01    rootdg    online
c0b0t2d0s7  sliced   -         -         online invalid
c0b0t3d0s7  sliced   -         -         error
bash-3.1#
```

Initializing the Volume Manager



VTOC of root after encapsulation

slice 0: STAND	permissions: VALID	starting sector: 48132 (cyl 2)	length: 81918 (5.10 cyls)
slice 1: SWAP	permissions: VALID UNMOUNTABLE	starting sector: 130050(cyl 8)	length: 864573 (53.82 cyls)
slice 2: ROOT	permissions: VALID	starting sector: 994623(cyl 61)	length: 7921446 (493.09 cyls)
slice 7: DISK	permissions: VALID UNMOUNTABLE	starting sector: 16065 (cyl 1)	length: 8900010 (554.00 cyls)
slice 13: BOOT	permissions: VALID UNMOUNTABLE	starting sector: 16065 (cyl 1)	length: 16002 (1.00 cyls)
slice 14: VOLPUBLIC	permissions: VALID UNMOUNTABLE	starting sector: 48132 (cyl 2)	length: 8867943 (552.00 cyls)
slice 15: VOLPRIVATE	permissions: VALID UNMOUNTABLE	starting sector: 32067 (cyl 1)	length: 16065 (1.00 cyls)

Before:

slice 0: STAND	permissions: VALID	starting sector: 32067 (cyl 1)	length: 81918 (5.10 cyls)
slice 1: SWAP	permissions: VALID UNMOUNTABLE	starting sector: 113985 (cyl 7)	length: 880638 (54.82 cyls)
slice 2: ROOT	permissions: VALID	starting sector: 994623 (cyl 61)	length: 7921446 (493.09 cyls)
slice 7: DISK	permissions: VALID UNMOUNTABLE	starting sector: 16065 (cyl 1)	length: 8900010 (554.00 cyls)
slice 13: BOOT	permissions: VALID UNMOUNTABLE	starting sector: 16065 (cyl 1)	length: 16002 (1.00 cyls)



Volume Manager Daemons

The Volume Manager Daemons are started automatically by the system start-up scripts. Presenting this information so you know they exist:

- **Vxconfigd**
 - Volume Manager Configuration daemon maintains volume manager disk and disk group configurations.
 - Communicates configuration changes to the kernel and modifies configuration information stored on the disk.
 - vxconfigd runs as a background process and can be running but not enabled. Use vxctl to check state of vxconfigd:

```
# vxctl mode  
mode enabled      :Indicates vxconfigd is both running and enabled.
```



Volume Manager Daemons, cont.

- Vxiod
 - Volume extended I/O daemon(s) allow for extended operations without blocking calling processes.
 - Kernel thread and is not visible to users via `ps`, must use the `vxiod` command to see the number of volume I/O daemons running.
 - May have several daemons running but usually one per processor, max 64.
 - May effect performance depending on operation.

`vxiod`

2 volume I/O daemons running

`vxiod` set 10

`vxiod`

10 volume I/O daemons running



Mirror root (boot) drive



Volume Manager Rootability

- Volume manager has capability of placing the *root*, *stand*, and *initial swap* under Volume Manager control – this is called *rootability*.
- The root (boot) disk can be put under VxVM control with a process of encapsulation, which converts existing partitions (slices) on the disk into volumes.
 - Swap area is referred to as *swap volume* and named *swapvol*
 - Root area is referred to as *root volume* and named *rootvol*
 - Stand area is referred to as *stand volume* and named *standvol*
- It is possible to mirror the *rootvol*, *swapvol*, and *standvol* to provide redundancy in the event of a disk failure.
- Note documentation for boot-time volume restrictions.
 - Reference VxVM Administrator's Guide pg 26 & 27.

HOWTO Mirror Root Drive



7. Mirror Boot drive

- CLI Operations VxVM User's Guide "Mirroring the Boot Disk"

```
# /etc/vx/bin/vxrootmir disk01
```

~ OR ~

HOWTO Mirror Root Drive



Menu Interface Operations, VxVM User's Guide "Mirroring Volumes on a VM Disk" #vxdiskadm

```
192.168.0.9 - PuTTY
Volume Manager Support Operations
Menu: VolumeManager/Disk

1      Add or initialize one or more disks
2      Encapsulate one or more disks
3      Remove a disk
4      Remove a disk for replacement
5      Replace a failed or removed disk
6      Mirror volumes on a disk
7      Move volumes from a disk
8      Enable access to (import) a disk group
9      Remove access to (deport) a disk group
10     Enable (online) a disk device
11     Disable (offline) a disk device
12     Mark a disk as a spare for a disk group
13     Turn off the spare flag on a disk
list   List disk information

?      Display help about menu
??     Display help about the menuing system
q      Exit from menus

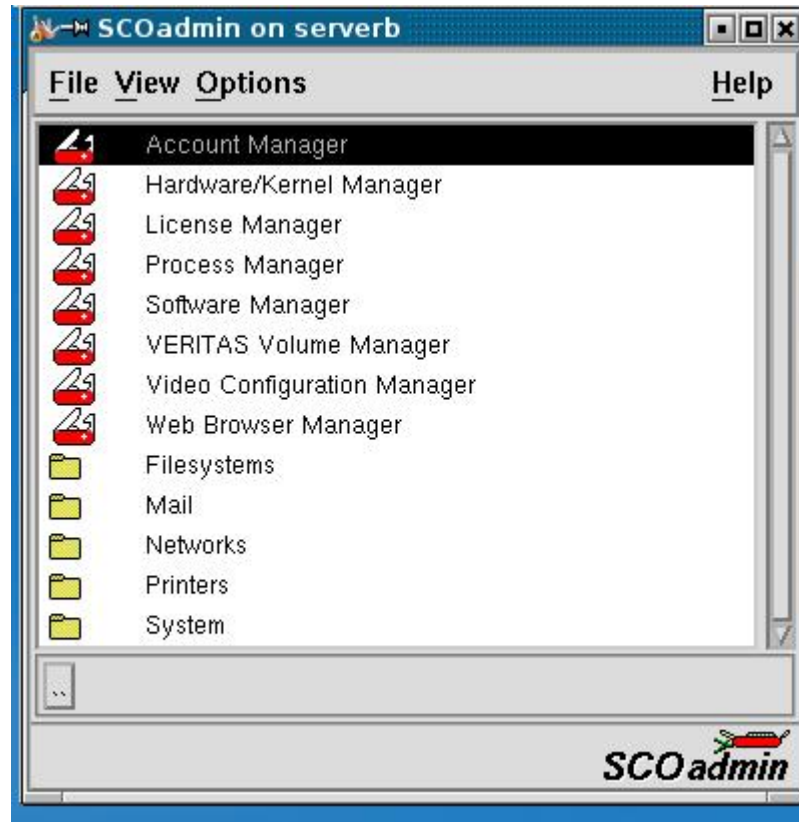
Select an operation to perform: █
```

HOWTO Mirror Root Drive



Visual Administrator

scoadmin



HOWTO Mirror Root Drive



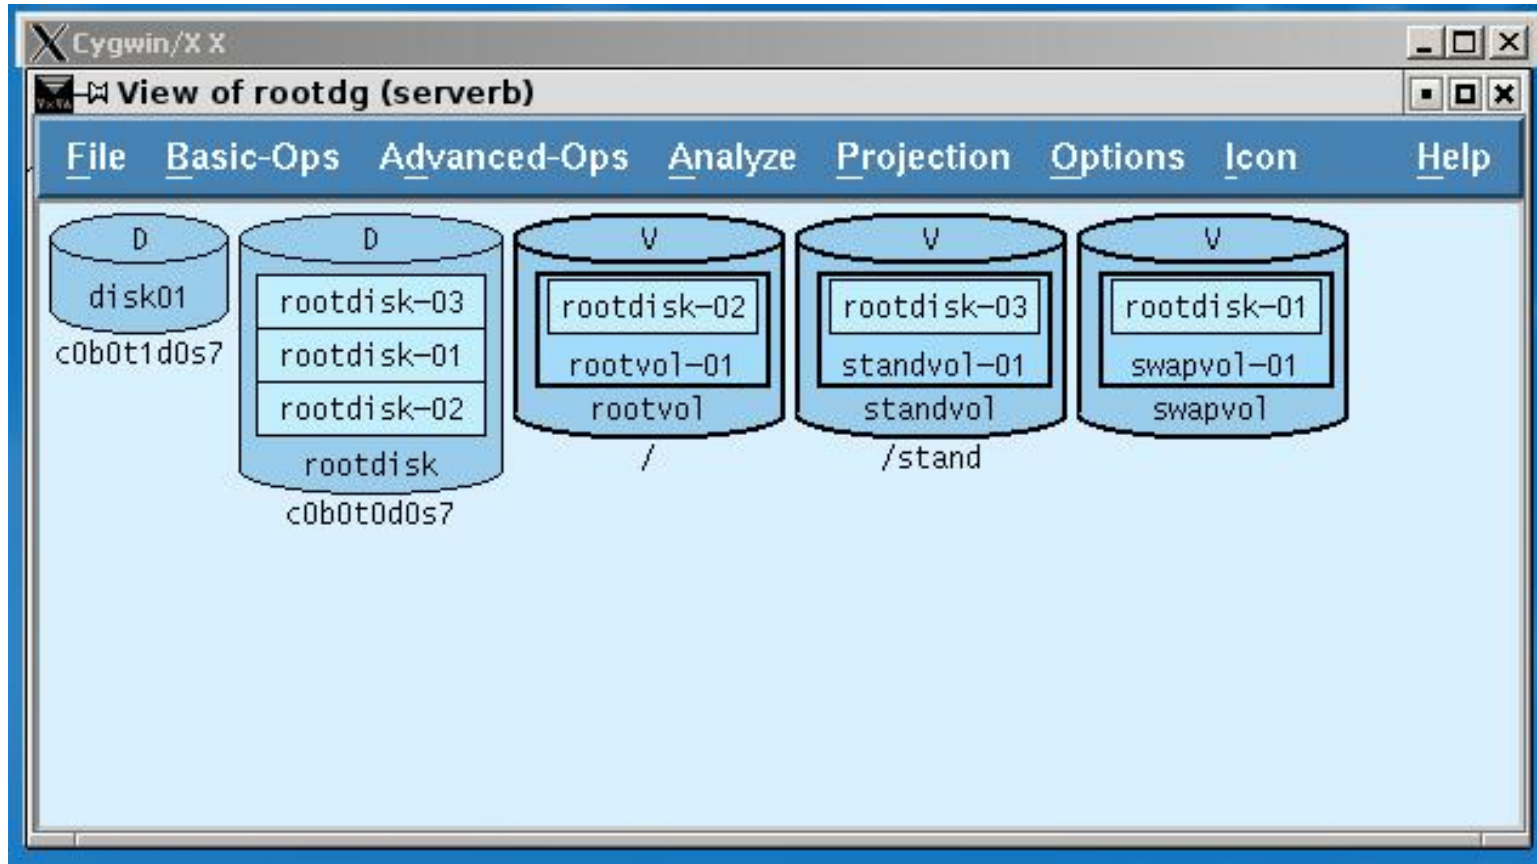
VxVM VA main screen:



HOWTO Mirror Root Drive



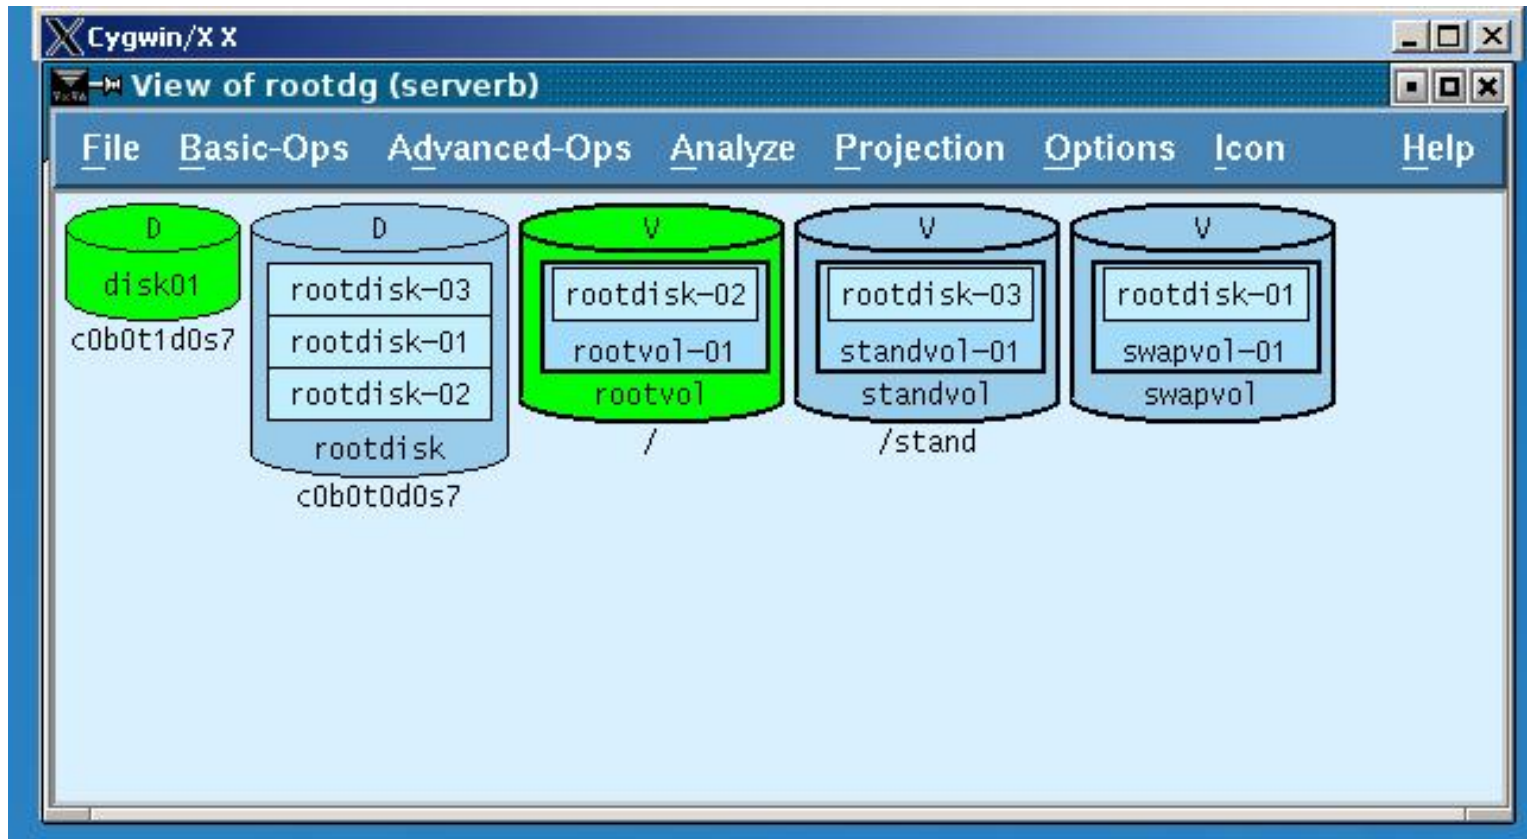
rootdg view



HOWTO Mirror Root Drive



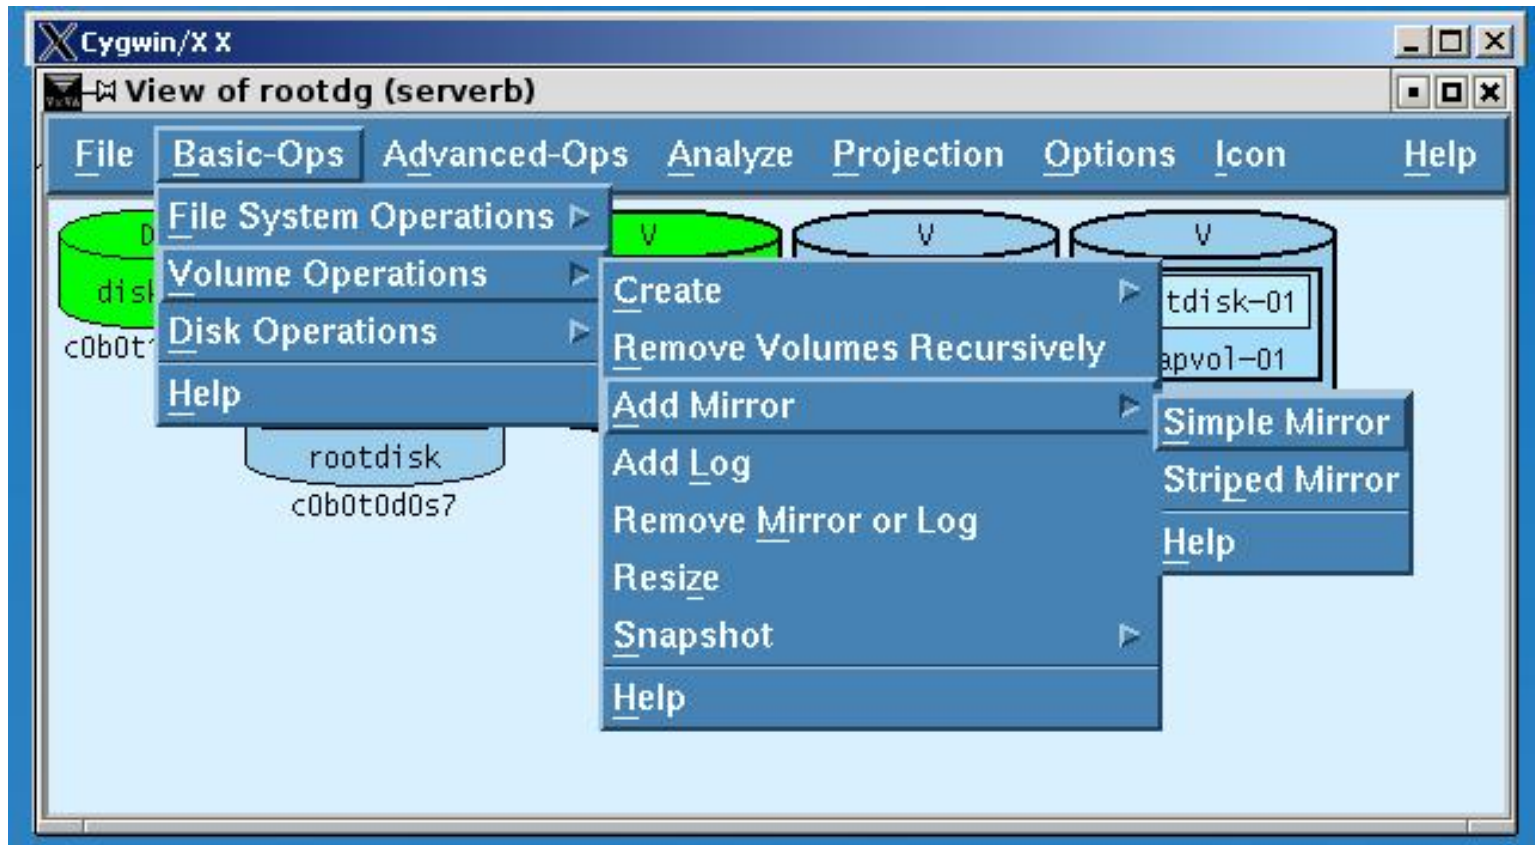
Select master volume (rootvol) and target disk



HOWTO Mirror Root Drive



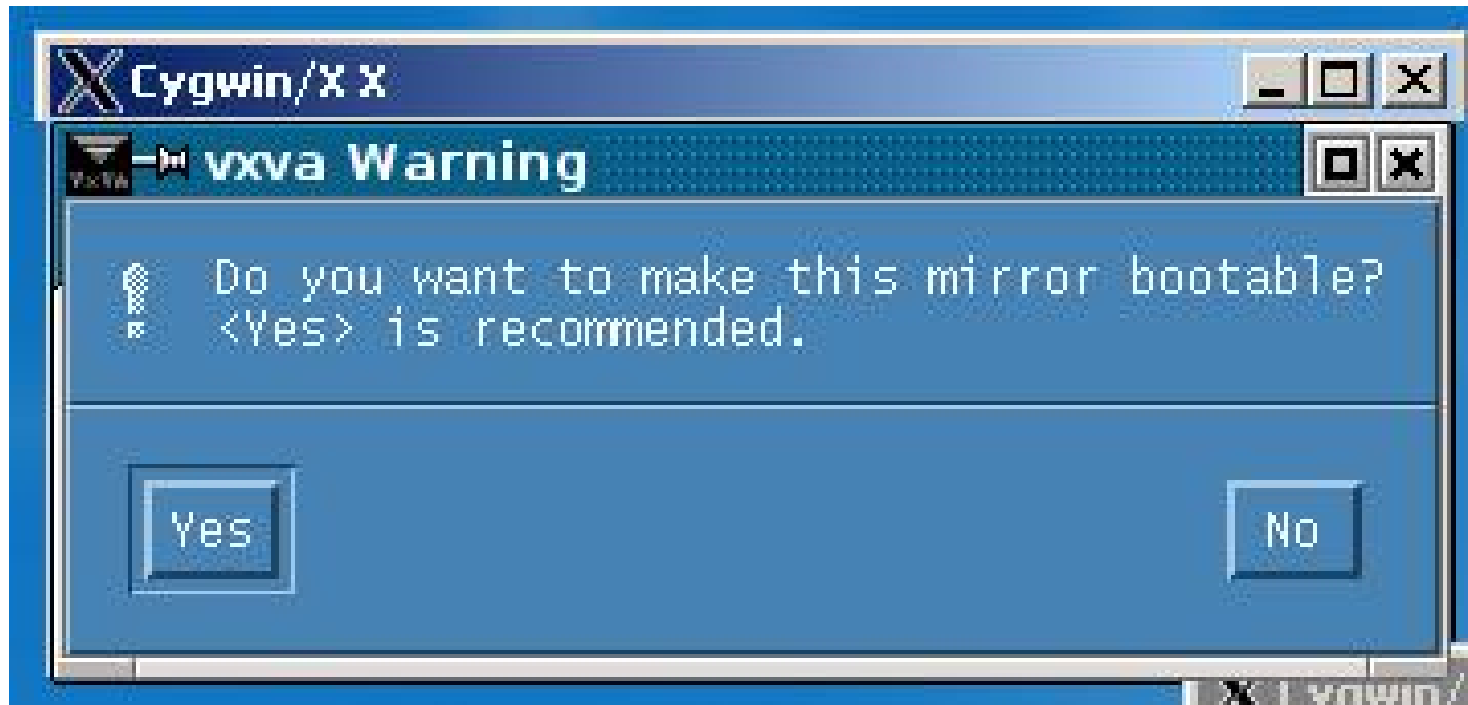
Basic-Ops -> Volume Operations -> Add Mirror -> Simple Mirror



HOWTO Mirror Root Drive



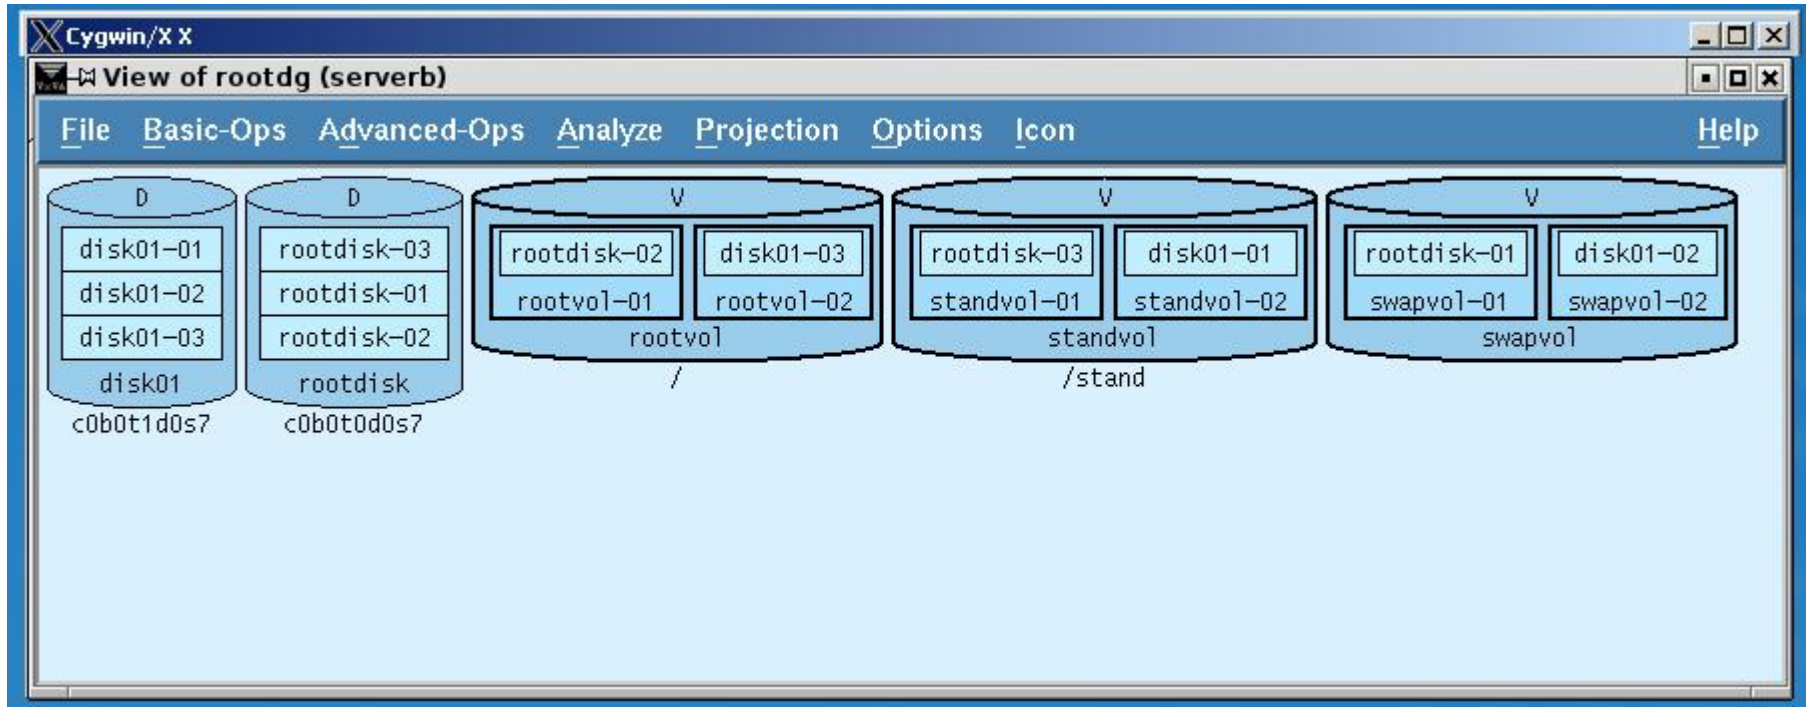
- Verification screen



HOWTO Mirror Root Drive



View of rootdg following mirror operation





Recovery



Recovery

- **Must read!** Appendix B, VxVM System Administration Guide
- Topics
 - VxVM UNIX boot process
 - Disk controller specifics
 - Configuring the system
 - Booting after failures
 - Creating and using VxVM Boot Floppy
 - Hot-Relocation and Boot Disk failures and recovery procedures
 - Re-Adding and replacing Boot disks
 - Plex and Volume States
 - Recovering the Volume Manager configuration



Reference VxVM Administrator's Guide, Appendix B

“Booting After Failures” – Many types of failures can prevent a system from booting.

- Try to identify the failure by messages left on the screen and correct if possible.
 - Visual inspections
 - Error logs
- Check the “State” of plex and volume.
- If the failure is not repairable, boot the system from an alternate boot disk.



Check the "State" of plex and volume - *vxprint*

Output of functioning system:

```
192.168.0.9 - PuTTY
bash-3.1# vxprint
Disk group: rootdg

TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg rootdg        rootdg         -        -        -        -        -        -
dm disk01        c0b0t1d0s7    -        8882922 -        -        -        -
dm rootdisk      c0b0t0d0s7    -        8867943 -        -        -        -

v  rootvol        root           ENABLED  7921446 -        ACTIVE  -        -
pl rootvol-01     rootvol        ENABLED  7921446 -        ACTIVE  -        -
sd rootdisk-02    rootvol-01     ENABLED  7921446 0        -        -        -
pl rootvol-02     rootvol        ENABLED  7921446 -        ACTIVE  -        -
sd disk01-03      rootvol-02     ENABLED  7921446 0        -        -        -

v  standvol       fsgen          ENABLED  81918   -        ACTIVE  -        -
pl standvol-01    standvol        ENABLED  81918   -        ACTIVE  -        -
sd rootdisk-03    standvol-01     ENABLED  81918   0        -        -        -
pl standvol-02    standvol        ENABLED  81918   -        ACTIVE  -        -
sd disk01-01      standvol-02     ENABLED  81918   0        -        -        -

v  swapvol        swap           ENABLED  864573  -        ACTIVE  -        -
pl swapvol-01     swapvol        ENABLED  864573  -        ACTIVE  -        -
sd rootdisk-01    swapvol-01     ENABLED  864573  0        -        -        -
pl swapvol-02     swapvol        ENABLED  864573  -        ACTIVE  -        -
sd disk01-02      swapvol-02     ENABLED  864573  0        -        -        -
bash-3.1# █
```

Recovery



/var/vxvm/vxconfigd.log

```
07/19 01:17:32: NOTICE: Offlining config copy 1 on disk c0b0t1d0s7:  
    Reason: Disk write failure  
07/19 01:17:32: NOTICE:  Detached disk disk01  
07/19 01:17:32: NOTICE:  Detached plex standvol-02 in volume standvol  
07/19 01:17:32: NOTICE:  Detached plex swapvol-02 in volume swapvol  
07/19 01:17:32: NOTICE:  Detached plex rootvol-02 in volume rootvol  
/var/vxvm/vxconfigd.log (END)
```



vxprint following failure – system still operational

```
192.168.0.9 - PuTTY
bash-3.1# vxprint
Disk group: rootdg

TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg rootdg        rootdg         -        -        -        -        -        -
dm disk01        -              -        -        -        NODEVICE -        -
dm rootdisk      c0b0t0d0s7    -        8867943 -        -        -        -

v  rootvol       root           ENABLED  7921446 -        ACTIVE  -        -
pl rootvol-01    rootvol       ENABLED  7921446 -        ACTIVE  -        -
sd rootdisk-02  rootvol-01    ENABLED  7921446 0        -        -        -
pl rootvol-02    rootvol       DISABLED 7921446 -        NODEVICE -        -
sd disk01-03    rootvol-02    DISABLED 7921446 0        NODEVICE -        -

v  standvol      fsgen         ENABLED  81918   -        ACTIVE  -        -
pl standvol-01  standvol      ENABLED  81918   -        ACTIVE  -        -
sd rootdisk-03 standvol-01    ENABLED  81918   0        -        -        -
pl standvol-02  standvol      DISABLED 81918   -        NODEVICE -        -
sd disk01-01    standvol-02    DISABLED 81918   0        NODEVICE -        -

v  swapvol       swap          ENABLED  864573  -        ACTIVE  -        -
pl swapvol-01   swapvol       ENABLED  864573  -        ACTIVE  -        -
sd rootdisk-01 swapvol-01    ENABLED  864573  0        -        -        -
pl swapvol-02   swapvol       DISABLED 864573  -        NODEVICE -        -
sd disk01-02    swapvol-02    DISABLED 864573  0        NODEVICE -        -
bash-3.1#
```

Recovery



vxdiskadm

```
192.168.0.9 - PuTTY
Volume Manager Support Operations
Menu: VolumeManager/Disk

1      Add or initialize one or more disks
2      Encapsulate one or more disks
3      Remove a disk
4      Remove a disk for replacement
5      Replace a failed or removed disk
6      Mirror volumes on a disk
7      Move volumes from a disk
8      Enable access to (import) a disk group
9      Remove access to (deport) a disk group
10     Enable (online) a disk device
11     Disable (offline) a disk device
12     Mark a disk as a spare for a disk group
13     Turn off the spare flag on a disk
list   List disk information

?      Display help about menu
??     Display help about the menuing system
q      Exit from menus

Select an operation to perform: █
```



Select option 4 – *Remove a disk for replacement*

```
192.168.0.9 - PuTTY
Remove a disk for replacement
Menu: VolumeManager/Disk/RemoveForReplace

Use this menu operation to remove a physical disk from a disk
group, while retaining the disk name. This changes the state
for the disk name to a "removed" disk. If there are any
initialized disks that are not part of a disk group, you will be
given the option of using one of these disks as a replacement.

Enter disk name [<disk>,list,q,?] disk01
```


Recovery



```
192.168.0.9 - PuTTY
Remove a disk for replacement
Menu: VolumeManager/Disk/RemoveForReplace

Use this menu operation to remove a physical disk from a disk
group, while retaining the disk name. This changes the state
for the disk name to a "removed" disk. If there are any
initialized disks that are not part of a disk group, you will be
given the option of using one of these disks as a replacement.

Enter disk name [<disk>,list,q,?] disk01

The following volumes will lose mirrors as a result of this
operation:

    rootvol standvol swapvol

No data on these volumes will be lost.
vxvm:vxdisk: ERROR: Device c0b0t1d0s7: get_contents failed:
    Disk device is offline

The requested operation is to remove disk disk01 from disk group
rootdg. The disk name will be kept, along with any volumes using
the disk, allowing replacement of the disk.

Select "Replace a failed or removed disk" from the main menu
when you wish to replace the disk.

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

Recovery



```
192.168.0.9 - PuTTY

Use this menu operation to remove a physical disk from a disk
group, while retaining the disk name. This changes the state
for the disk name to a "removed" disk. If there are any
initialized disks that are not part of a disk group, you will be
given the option of using one of these disks as a replacement.

Enter disk name [<disk>,list,q,?] disk01

The following volumes will lose mirrors as a result of this
operation:

    rootvol standvol swapvol

No data on these volumes will be lost.
vxvm:vxdisk: ERROR: Device c0b0t1d0s7: get_contents failed:
    Disk device is offline

The requested operation is to remove disk disk01 from disk group
rootdg. The disk name will be kept, along with any volumes using
the disk, allowing replacement of the disk.

Select "Replace a failed or removed disk" from the main menu
when you wish to replace the disk.

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

Removal of disk disk01 completed successfully.

Remove another disk? [y,n,q,?] (default: n)
```



`vxprint state` following `vxdiskadm` removal:

```
192.168.0.9 - PuTTY
bash-3.1# vxprint
Disk group: rootdg

TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg rootdg        rootdg         -        -        -        -        -        -
dm disk01        -              -        -        -        REMOVED -        -
dm rootdisk      c0b0t0d0s7    -        8867943 -        -        -        -
v  rootvol       root           ENABLED  7921446 -        ACTIVE  -        -
pl rootvol-01    rootvol       ENABLED  7921446 -        ACTIVE  -        -
sd rootdisk-02  rootvol-01    ENABLED  7921446 0        -        -        -
pl rootvol-02    rootvol       DISABLED 7921446 -        REMOVED -        -
sd disk01-03    rootvol-02    DISABLED 7921446 0        REMOVED -        -
v  standvol      fsgen         ENABLED  81918   -        ACTIVE  -        -
pl standvol-01  standvol      ENABLED  81918   -        ACTIVE  -        -
sd rootdisk-03  standvol-01  ENABLED  81918   0        -        -        -
pl standvol-02  standvol      DISABLED 81918   -        REMOVED -        -
sd disk01-01    standvol-02  DISABLED 81918   0        REMOVED -        -
v  swapvol       swap          ENABLED  864573  -        ACTIVE  -        -
pl swapvol-01  swapvol       ENABLED  864573  -        ACTIVE  -        -
sd rootdisk-01  swapvol-01    ENABLED  864573  0        -        -        -
pl swapvol-02  swapvol       DISABLED 864573  -        REMOVED -        -
sd disk01-02    swapvol-02    DISABLED 864573  0        REMOVED -        -
bash-3.1#
```



Power down system and replace failed disk

Recovery



mkdev hd

Write master boot area.

Initialize if needed.



Vxdiskadm – replace failed disk

```
192.168.0.9 - PuTTY
Replace a failed or removed disk
Menu: VolumeManager/Disk/ReplaceDisk

Use this menu operation to specify a replacement disk for a disk
that you removed with the "Remove a disk for replacement" menu
operation, or that failed during use. You will be prompted for
a disk name to replace and a disk device to use as a replacement.
You can choose an uninitialized disk, in which case the disk will
be initialized, or you can choose a disk that you have already
initialized using the Add or initialize a disk menu operation.

Select a removed or failed disk [<disk>,list,q,?] disk01
```

Recovery



```
192.168.0.9 - PuTTY

Replace a failed or removed disk
Menu: VolumeManager/Disk/ReplaceDisk

Use this menu operation to specify a replacement disk for a disk
that you removed with the "Remove a disk for replacement" menu
operation, or that failed during use. You will be prompted for
a disk name to replace and a disk device to use as a replacement.
You can choose an uninitialized disk, in which case the disk will
be initialized, or you can choose a disk that you have already
initialized using the Add or initialize a disk menu operation.

Select a removed or failed disk [<disk>,list,q,?] disk01

Select disk device to initialize [<address>,list,q,?] list

DEVICE      DISK      GROUP      STATUS
c0b0t0d0    rootdisk  rootdg     online
c0b0t1d0    -         -          online
c0b0t2d0    -         -          online invalid
c0b0t3d0    -         -          error

Select disk device to initialize [<address>,list,q,?] c0b0t1d0
```

Recovery



```
192.168.0.9 - PuTTY
Replace a failed or removed disk
Menu: VolumeManager/Disk/ReplaceDisk

Use this menu operation to specify a replacement disk for a disk
that you removed with the "Remove a disk for replacement" menu
operation, or that failed during use. You will be prompted for
a disk name to replace and a disk device to use as a replacement.
You can choose an uninitialized disk, in which case the disk will
be initialized, or you can choose a disk that you have already
initialized using the Add or initialize a disk menu operation.

Select a removed or failed disk [<disk>,list,q,?] disk01

Select disk device to initialize [<address>,list,q,?] list

DEVICE      DISK      GROUP      STATUS
c0b0t0d0    rootdisk  rootdg     online
c0b0t1d0    -         -          online
c0b0t2d0    -         -          online invalid
c0b0t3d0    -         -          error

Select disk device to initialize [<address>,list,q,?] c0b0t1d0

The following disk device appears to have been initialized already.
The disk is listed as added to a disk group, but is not currently
active. Output format: [Device_Name,Disk_Access_Name,Disk_Group]

[c0b0t1d0,c0b0t1d0s7,rootdg]

Use this device? [y,n,q,?] (default: y) █
```




Successful replacement

```
192.168.0.9 - PuTTY

Select disk device to initialize [<address>,list,q,?] c0b0t1d0

The following disk device appears to have been initialized already.
The disk is listed as added to a disk group, but is not currently
active.  Output format: [Device_Name,Disk_Access_Name,Disk_Group]

[c0b0t1d0,c0b0t1d0s7,rootdg]

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

The following disk you selected for use appears to already have
been initialized for the Volume Manager.  If you are certain the
disk has already been initialized for the Volume Manager, then you
do not need to reinitialize the disk device.
Output format: [Device_Name]

c0b0t1d0

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

The requested operation is to use the initialized device c0b0t1d0
to replace the removed or failed disk disk01 in disk group rootdg.

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

Replacement of disk disk01 in group rootdg with disk device
c0b0t1d0 completed successfully.

Replace another disk? [y,n,q,?] (default: n) █
```

Recovery



vxprint showing recovery in progress:

```
192.168.0.9 - PuTTY
bash-3.1# vxprint
Disk group: rootdg

TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg rootdg        rootdg         -        -        -        -        -        -
dm disk01        c0b0t1d0s7    -        8882922 -        -        -        -
dm rootdisk      c0b0t0d0s7    -        8867943 -        -        -        -

v rootvol        root           ENABLED  7921446 -        ACTIVE  ATT1    -
pl rootvol-01    rootvol        ENABLED  7921446 -        ACTIVE  -        -
sd rootdisk-02   rootvol-01     ENABLED  7921446 0        -        -        -
pl rootvol-02    rootvol        ENABLED  7921446 -        STALE   ATT     -
sd disk01-03     rootvol-02     ENABLED  7921446 0        -        -        -

v standvol       fsgen          ENABLED  81918   -        ACTIVE  -        -
pl standvol-01   standvol        ENABLED  81918   -        ACTIVE  -        -
sd rootdisk-03   standvol-01     ENABLED  81918   0        -        -        -
pl standvol-02   standvol        DISABLED 81918   -        RECOVER -        -
sd disk01-01     standvol-02     ENABLED  81918   0        -        -        -

v swapvol        swap           ENABLED  864573  -        ACTIVE  -        -
pl swapvol-01    swapvol        ENABLED  864573  -        ACTIVE  -        -
sd rootdisk-01   swapvol-01     ENABLED  864573  0        -        -        -
pl swapvol-02    swapvol        DISABLED 864573  -        RECOVER -        -
sd disk01-02     swapvol-02     ENABLED  864573  0        -        -        -
bash-3.1# █
```



vxprint status after mirroring replacement disk:

```
192.168.0.9 - PuTTY
bash-3.1# vxprint
Disk group: rootdg

TY NAME          ASSOC          KSTATE  LENGTH  PLOFFS  STATE  TUTILO  PUTILO
dg rootdg        rootdg         -        -        -        -        -        -
dm disk01        c0b0t1d0s7    -        8882922 -        -        -        -
dm rootdisk      c0b0t0d0s7    -        8867943 -        -        -        -

v  rootvol        root           ENABLED  7921446 -        ACTIVE  -        -
pl rootvol-01     rootvol        ENABLED  7921446 -        ACTIVE  -        -
sd rootdisk-02    rootvol-01     ENABLED  7921446 0        -        -        -
pl rootvol-02     rootvol        ENABLED  7921446 -        ACTIVE  -        -
sd disk01-03      rootvol-02     ENABLED  7921446 0        -        -        -

v  standvol       fsgen          ENABLED  81918   -        ACTIVE  -        -
pl standvol-01    standvol       ENABLED  81918   -        ACTIVE  -        -
sd rootdisk-03    standvol-01    ENABLED  81918   0        -        -        -
pl standvol-02    standvol       ENABLED  81918   -        ACTIVE  -        -
sd disk01-01      standvol-02    ENABLED  81918   0        -        -        -

v  swapvol        swap           ENABLED  864573  -        ACTIVE  -        -
pl swapvol-01     swapvol        ENABLED  864573  -        ACTIVE  -        -
sd rootdisk-01    swapvol-01     ENABLED  864573  0        -        -        -
pl swapvol-02     swapvol        ENABLED  864573  -        ACTIVE  -        -
sd disk01-02      swapvol-02     ENABLED  864573  0        -        -        -
bash-3.1#
```



Online Data Manager (ODM) summary

- Cost effective enterprise solution to protect and manage data.
- Provides software RAID Levels 0, 1, 5, 10.
- On-line disk performance analysis.
- It includes a graphical Visual Administrator for exceptional ease-of-use.
- Supports online filesystem resizing.
- Boot disk data protection with mirroring.



Q & A