

VERITAS Cluster Server™

Application for SunStoreEDGE A5000 Disk Array

Solaris

October 2000
30-000072-399


VERITAS

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 © 2000 VERITAS Software Corporation. All rights reserved. VERITAS is a registered trademark of VERITAS Software Corporation in the US and other countries. The VERITAS logo and VERITAS Cluster Server are trademarks of VERITAS Software Corporation. All other trademarks or registered trademarks are the property of their respective owners.

Printed in the USA, October 2000.

VERITAS Software Corporation
1600 Plymouth St.
Mountain View, CA 94043
Phone 650-335-8000
Fax 650-335-8050
www.veritas.com



Contents

Installing and Configuring the Sun StorEdge A5000 Disk Array	1
Introduction	1
Hardware Requirements	1
SBus Systems	1
PCI Bus Systems	1
Software Requirements	2
Recommended Patches	2
Solaris 2.5.1 8/97 With Maintenance Update (MU)	2
Solaris 2.6 With MU	3
Configuration Modes	4
About Dynamic Multipathing	5
Setting Up the Hardware	6
Creating Disk Groups and Volumes	9
Dynamic Multipathing with VCS	10
NFS Failover Setup	10
Remajoring Devices	11
Determining if Your Shared Device Needs Remajoring	11
Remajoring Your Shared Device	11
Glossary	13





Preface

This document provides instructions on how to install and configure the Sun StorEDGE A5000 disk array for the VERITAS® Cluster Server™ (VCS) failover environment. For information about VCS, refer to the *VERITAS Cluster Server User's Guide*. For more information about the Sun StorEDGE A5000 disk array, refer to the appropriate Sun StorEDGE A5000 documentation.

Technical Support

For assistance with this product, or information regarding VERITAS service packages, contact Technical Support at the numbers listed below. You may also contact Technical Support via email at support@veritas.com.

VCS on UNIX: from U.S and Canada, call 800.342.0652.

For Customers Outside U.S. and Canada

From Europe, the Middle East, or Asia, visit the Technical Support website at <http://support.veritas.com> for a list of each country's contact information.



Conventions

Typeface	Usage
<code>courier</code>	computer output, command references within text
<code>courier</code> (bold)	user input and commands, keywords in grammar syntax
<i>italic</i>	new terms, book titles, emphasis
<i>italic</i>	variables within a command
Symbol	Usage
#	UNIX superuser prompt (for all shells)
%	C shell prompt
\$	Bourne/Korn shell prompt



Installing and Configuring the Sun StoreEDGE A5000 Disk Array

Introduction

VCS operates with the Sun StoreEDGE A5000 disk array without additional software. This document instructs you on how to configure the Sun StoreEDGE A5000 for use with VCS.

Hardware Requirements

SBus Systems

- ◆ One or more Sun StoreEDGE A5000 or A5100 disk arrays.
- ◆ Two Sun Enterprise E3000, E4000, E5000, E6000, or E10000 servers.
- ◆ Two FC100/SBus host bus adapters. Each FC100/S adapter has two fibre-channel ports. Although one card per server is required, two cards prevent a single-card point of failure.

PCI Bus Systems

- ◆ Two Sun Enterprise 450 servers (see page 3).
- ◆ Two FC100/PCI host bus adapters. Install and configure the SUNWifp and SUNWifvp packages from the CD included with the FC100 cards.



Software Requirements

- ◆ Solaris 2.6, or 2.5.1 8/97 with maintenance update.
- ◆ VCS 1.0.2 or later.
- ◆ Sun EVM (shipped with Sun StorEDGE A5000) or VERITAS Volume Manager™ (VxVM), version 2.5.1 or later.

Note For E450 support, only VxVM 3.0.1 or later and Solaris 2.6 are supported.

Recommended Patches

The following patches are currently available at Sun's SunSolve website. We recommend you use the most recent releases of these patches. Patch updates are also available at this site.

Solaris 2.5.1 8/97 With Maintenance Update (MU)

The following patches are required to run the A5000 disk array under Solaris 2.5.1 8/97 with MU:

Patch Number	Description
104708-12	SunOS 5.5.1: <code>ssd</code> , <code>pln</code> , <code>soc</code> , <code>ssaadm</code> and <code>ssafirmware</code> patch
105298-02	SunOS 5.5.1: AP 2.0: <code>apdisk/ap_daemon/mqfe</code> fixes
105324-01	SunOS 5.5.1: <code>ses</code> driver patch
105310-05	SunOS 5.5.1: <code>socal</code> and <code>sf</code> driver patches, <code>luxadm</code> utility patches

- ✓ Install the four patches on both hosts in the order listed in the table above.
- ✓ After installing the last patch on both hosts, follow the special install instructions from one host to upgrade the A5000 firmware to version 1.05, and FC100/S fcode level to version 1.11.

The following patch upgrades ST19171FC disk firmware to level 1178. Run this patch from one host:

Patch Number	Description
106129-02	SunOS 5.5.1, 5.6: hardware, 9GB disks. Download program and firmware 0411 for ST19171FC.



Solaris 2.6 With MU

The following patches are required to run the A5000 disk array under Solaris 2.6 with MU:

Patch	Description
105356-04	SunOS 5.6: /kernel/drv/ssd patch
105357-01	SunOS 5.6: /kernel/drv/ses patch
105375-04	SunOS 5.6: sf and soca1 driver patch

- ✓ Install the three patches on both hosts in the order listed in the table above.
- ✓ After installing the last patch on both hosts, follow the special install instructions from one host to upgrade A5000 firmware to version 1.05, and FC100/S fcode level to version 1.11.

The following patch upgrades ST19171FC disk firmware to level 1178. Run this patch from one host:

Patch	Description
106129-02	SunOS 5.6, 5.6: hardware, 9GB disks. Download program and firmware 0411 for ST19171FC.

For E450 Servers

The following patch is required to run the A5000 disk array on E450 servers:

Patch	Description
107280-01	SunOS 5.6 /kernel/drv/ifp patch

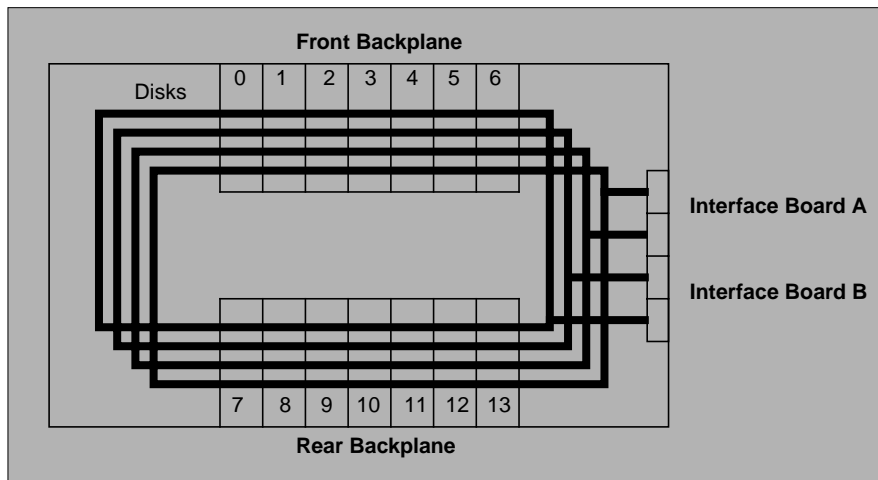


Configuration Modes

The Sun StorEDGE A5000 disk array contains fourteen drive bays: seven are located along the front of the disk array, and seven along the rear. The A5000 logically divides the disks in the drive bays between the two backplanes that connect the front and rear sets of disks. You can configure the A5000 in two modes: single-loop and split-loop. *For the greatest degree of high availability, we recommend you configure the A5000 in single-loop mode.*

Single-Loop Mode

Both backplanes are interconnected to create a single loop. This enables each gigabit interface converter (GBIC) to recognize all disks contained in the disk array.



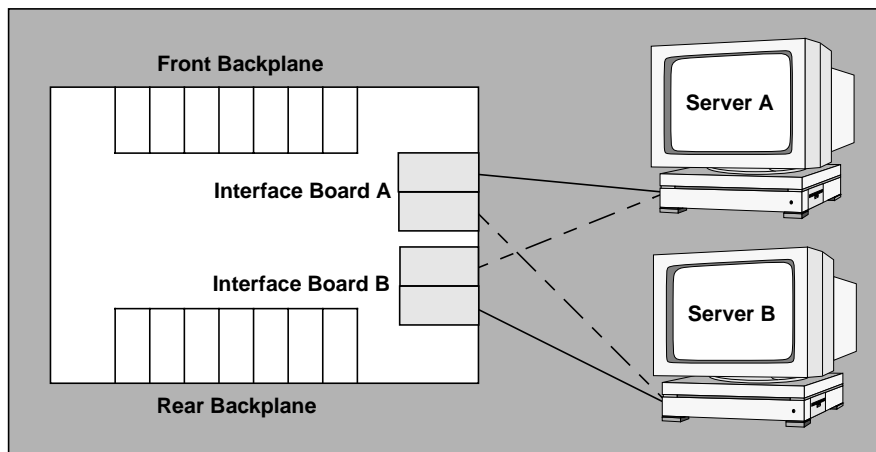
A5000 Single-Loop Configuration

Split-Loop Mode

The backplanes remain separated. This enables one GBIC on each interface board to recognize a loop containing only front backplane disks, and the other GBIC on each interface board to recognize a loop containing only rear backplane disks.

About Dynamic Multipathing

Dynamic multipathing (DMP) provided by VxVM transparently recovers from failures caused by faulty host adapters, cables, or interface boards. The VxVM software was designed to notice a controller failure, and implement the appropriate recovery procedures by rerouting requests to the volume through an alternate path. Because DMP is an essential component of the A5000's redundancy, we recommend each host have two fibre-optic connections to the disk array. For information on DMP with VCS, see [“Dynamic Multipathing with VCS”](#) on page 10.




A5000 Dual-Host Configuration



Setting Up the Hardware

Perform the following steps on each system:

1. Configure VCS on all hosts as described in the *VERITAS Cluster Server Installation Guide*.
2. Verify that all hardware, software, and required patches are properly installed. If necessary, refer to page 1.
3. Use fibre-optic cables to connect each host server's FC100 adapter to the GBICs located on the A5000 interface boards.
4. Connect one adapter port from each host to interface board A and the other to interface board B. Connecting each host to each interface board enables DMP to handle an interface board failure (rear view shown below).



5. Set the disk array to operate in single-loop mode by using the front panel module (FPM) screen on the disk array. Refer to the *Sun StorEDGE A5000 Disk Array Installation and Service Manual* for instructions.

Note that you can daisy-chain up to four Sun StorEDGE A5000 disk arrays. Refer to the *Sun StorEDGE A5000 Disk Array Hardware Configuration Guide* for instructions.

6. Reboot your host system:

```
serverA # reboot -- -r
```

7. Before you configure volumes, verify both servers have rebooted by entering the `format` command. When the `format` command displays all A5000 drives under two controllers, the installation is confirmed:

```
serverA # format
```

Output resembles:

```
Searching for disks...done
```

```
AVAILABLE DISK SELECTIONS:
```

```
0. c0t0d0 <SUN2.1G cyl 2733 alt 2 hd 19 sec 80>
   /sbus@1f,0/espdma@e,8400000/esp@e,8800000/sd@0,0
1. clt64d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w21000020370499a6,0
2. clt65d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049be1,0
3. clt66d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w210000203704981f,0
4. clt67d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049cfc,0
5. clt68d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w21000020370494be,0
6. clt69d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049f02,0
7. clt70d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049a2d,0
8. clt80d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049d1e,0
9. clt81d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049d71,0
10. clt82d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w21000020370494b3,0
11. clt83d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049a6e,0
12. clt84d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049c7f,0
13. clt85d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049b33,0
14. clt86d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2100002037049bf8,0
15. c2t64d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w21000020370499a6,0
16. c2t65d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049be1,0
17. c2t66d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
   /sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w210000203704981f,0
18. c2t67d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
```



```
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049cfc,0
19. c2t68d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w21000020370494be,0
20. c2t69d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049f02,0
21. c2t70d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049a2d,0
22. c2t80d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049d1e,0
23. c2t81d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049d71,0
24. c2t82d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w21000020370494b3,0
25. c2t83d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049a6e,0
26. c2t84d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049c7f,0
27. c2t85d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049b33,0
28. c2t86d0 <SUN9.0G cyl 4924 alt 2 hd 27 sec 133>
/sbus@1f,0/SUNW,socal@1,0/sf@1,0/ssd@w2100002037049bf8,0
Specify disk (enter its number):
```



Creating Disk Groups and Volumes

In the example beginning on page 7, the `format` command shows that disks 1–14 under controller 1 are actually the same as disks 15–28. The target number and World Wide Name (WWN) number are identical between these device nodes. However, the `vxdisk list` command of the Volume Manager DMP feature shows that VxVM only allows the configuration under one controller, as illustrated below:

```
serverA # vxdisk list
```

DEVICE	TYPE	DISK	GROUP	STATUS
c0t0d0s2	sliced	rootdisk	rootdg	online
c1t64d0s2	sliced	-	-	online
c1t65d0s2	sliced	-	-	online
c1t66d0s2	sliced	-	-	online
c1t67d0s2	sliced	-	-	online
c1t68d0s2	sliced	-	-	online
c1t69d0s2	sliced	-	-	online
c1t70d0s2	sliced	-	-	online
c1t80d0s2	sliced	-	-	online
c1t81d0s2	sliced	-	-	online
c1t82d0s2	sliced	-	-	online
c1t83d0s2	sliced	-	-	online
c1t84d0s2	sliced	-	-	online
c1t85d0s2	sliced	-	-	online
c1t86d0s2	sliced	-	-	online

You can now create the disk groups and volumes using the disks listed under controller 1. For more information on creating disk groups and volumes, refer to the *VERITAS Volume Manager User's Guide*.



Dynamic Multipathing with VCS

If a host fails or if a VCS agent detects a failure, VCS fails over volume control to the remote host. If a failure occurs within a host adapter, fibre-optic cable, or A5000 interface board, DMP within VxVM detects the failure and transparently moves I/O control to the other operational loop (controller 2 in the previous example).

If a failure occurs that causes DMP to fail over disk control between controllers on a VCS server, you must run the `vxdtcl enable` command on the server when the failed path is restored. This command updates the Volume Manager configuration database, enabling DMP functionality for that device.

NFS Failover Setup

If you are using a shared file system for NFS failover, the major and minor numbers for the device corresponding to the shared file system must match on both servers.

To enable VCS to control NFS daemon start and stop, configure the NFS agent as specified in the *VERITAS Cluster Server Installation Guide*.



Remajoring Devices

Determining if Your Shared Device Needs Remajoring

The following example demonstrates how to determine if a major number must be changed. In this example, VCS serverA and serverB have a shared device `/dev/vx/dsk/dg1`.

1. List the device on serverA:

```
serverA # ls -lL /dev/vx/dsk/dg1
```

Output resembles:

```
total 0
brw----- 1 root root 87,45000 Dec 03 16:00 voll
```

2. List the device on serverB:

```
serverB # ls -lL /dev/vx/dsk/dg1
total 0
brw----- 1 root root 89,45000 Dec 03 16:00 voll
```

ServerA shows the shared device with the major number as 87 and serverB as 89. The number is not identical on both servers, so you must change the shared device to a major number that is unique to the servers, and identical between them.

Remajoring Your Shared Device

If your shared device requires remajoring, refer to the `haremajor` command to change the major number for the `vxio` driver.





Glossary

adapter

Fibre-channel host adapter located on the server.

backplane

Routes power, control, sense, and fibre-channel signals from the interconnect assembly to the A5000 disk array disk drive slots.

DMP

VERITAS Volume Manager uses dynamic multipathing (DMP) to detect an I/O path failure. DMP implements the appropriate recovery steps by transparently rerouting requests to a volume through an alternate path, without interrupting the end user.

FC-AL

Fibre-Channel Arbitrated-Loop topology that enables you to attach multiple devices in a loop without hubs and switches.

FPM

The front panel module (FPM) uses an electroluminescent display panel to display A5000 configuration, status, and diagnostic information. The FPM touch-screen enables you to monitor and change configurations.

GBIC

The gigabit interface converter (GBIC) is a hot-pluggable unit that converts any of the standard fibre-channel connectors and signaling technologies. The A5000 disk array contains four GBIC adapters, two on each interface board.

interface board

The interface board provides intelligent controls in the A5000 enclosure. It senses and sets required environmental service signals by conditions within the enclosure. The interface board receives, relocks, and provides bypass services for two independent FC-AL ports. It also manages internal-loop configuration.



luxadm(1M)

A command line that the UNIX utility uses to control the A5000 disk array.

single-loop

An A5000 configuration setting that allows A5000 disks to be seen within a single FC-AL loop.

split-loop

An A5000 configuration setting in which the front backplane disks are seen within one FC-AL loop, and the rear backplane disks are seen within another FC-AL loop.

WWN

World Wide Name (WWN). A 16-character hexadecimal identification number burned into the EEPROM on each disk.

