

StorageWorks™ Solutions

Getting Started with the StorageWorks Network Storage Array

Order Number: EK-SWXNA-IG. A02

This guide describes the procedures necessary to install, configure, and use the StorageWorks Network Storage Array. Mechanical cabinet configuration and power configuration as well as software installation and configuration are covered in this guide.

September 1996

While Digital believes the information included in this publication is correct as of the date of publication, it is subject to change without notice.

Digital Equipment Corporation makes no representations that the interconnection of its products in the manner described in this document will not infringe existing or future patent rights, nor do the descriptions contained in this document imply the granting of licenses to make, use, or sell equipment or software in accordance with the description.

© Digital Equipment Corporation, 1996.
All Rights Reserved.
Printed in the United States of America.

NOTE: This equipment generates, uses, and may emit radio frequency energy. The equipment has been type tested and found to comply with the limits for a Class A digital device pursuant to Part 15 of the FCC rules. These limits are designed to provide reasonable protection against harmful interference in a residential installation.

Any changes or modifications made to this equipment may void the user's authority to operate the equipment.

Operation of this equipment in a residential area may cause interference in which case the user at his own expense will be required to take whatever measures may be required to correct the interference.

Alpha, Bookreader, Digital, DECnet, DECwindows, OpenVMS, POLYCENTER, StorageWorks, VMS, VMScluster, and the DIGITAL logo are trademarks of Digital Equipment Corporation.

Motif is a registered trademark of the Open Software Foundation, Inc.

PostScript is a registered trademark of Adobe Systems, Inc.

UNIX is a registered trademark in the United States and other countries licensed exclusively through X/Open Company Ltd.

This document was prepared using VAX DOCUMENT Version 2.1.

Contents

Preface	vii
Manufacturer's Declarations	xi
1 Introduction	
1.1 System Overview	1-1
1.1.1 SWXNA-Bx/Gx Description	1-1
1.1.2 SWXNA-Cx Description	1-2
1.1.3 SWXNA-Ax/Fx Description	1-3
1.1.4 SWXNA-Ex/Jx Description	1-4
1.1.5 SWXNA-DA/HA Description	1-5
1.1.6 SWXNA-xx Series Software and Firmware Components	1-5
1.2 Available Upgrades	1-6
1.2.1 StorageWorks Network Storage Array Specifications	1-7
1.3 Major Components	1-9
1.3.1 Server Processor	1-10
1.3.1.1 Shadowed System Disk	1-11
1.3.1.2 CD-ROM Drive	1-11
1.3.1.3 3.5-Inch Floppy Drive	1-13
1.3.1.4 Operator Control Panel	1-13
1.3.2 Device Channel Processor	1-14
1.3.2.1 Read/Write Cache	1-14
1.3.2.2 RAIDset Configuration Support	1-15
1.3.3 StorageWorks Storage Devices	1-15
1.4 Using the CD-ROM	1-15
1.4.1 Handling the CD-ROM	1-16
1.4.2 Accessing Documentation	1-16
2 Preparing Your Site	
2.1 Power Requirements	2-1
2.2 Grounding Requirements	2-2
2.3 Safety and Service Requirements	2-3
3 Unpacking, Placing, and Applying Power to the StorageWorks Network Storage Array	
3.1 Unpacking the Cabinet	3-1
3.2 Removing the Cabinet from the Pallet	3-4
3.3 Placing the Cabinet	3-7
3.4 Leveling the Cabinet	3-8
3.5 Installing the Skirt Kit	3-8
3.6 Inspecting the Cabinet	3-9

3.7	Connecting a Terminal to the StorageWorks Network Storage Array	3-11
3.8	Applying Power to the Cabinet	3-12
4	Customizing Network Array Operating System Parameters	
4.1	Software Configuration Overview	4-1
4.2	VMScluster System Considerations	4-2
4.2.1	Considering the VMScluster Quorum Scheme	4-2
4.2.2	Volume Shadowing Interoperability	4-3
4.3	Executing the Software Customization Procedure	4-3
5	Configuring the StorageWorks Network Storage Array Storage Devices	
5.1	General Considerations for Configuring a StorageWorks Storage System	5-1
5.2	Running CFMENU from the StorageWorks Network Storage Array Terminal	5-2
5.3	Configuring Storage Devices Using CFMENU	5-5
5.3.1	Using CFMENU	5-5
5.3.2	Considerations for Using CFMENU	5-5
5.4	Modifying Allocation Classes	5-5
5.4.1	Configuring Storage for Additional Device Channel Processor Pairs	5-7
5.5	Saving the Configuration Setup	5-7
6	Managing the Network Storage Array's Read Cache	
6.1	Using the PerfectCache Menu	6-2
6.1.1	Manage Cacher	6-2
6.1.1.1	Using the Cache Manager Screen	6-3
6.1.2	Display Caching Summary	6-4
6.1.3	Display Caching Summary Continuously	6-5
6.1.4	Reset Caching Statistics	6-5
6.2	PerfectCache Commands	6-5
7	Network File System	
7.1	Network File System	7-1
7.1.1	NFS Server	7-1
7.1.2	NFS Client	7-1
7.1.3	PC-NFS Client Daemon	7-2
7.2	Checklist for Configuring the NFS Server	7-2
7.3	NFS Configuration Session	7-7
7.4	Client Authorization and Served Disk Mapping Configuration	7-11
8	Integrating the Network Array into the Network System	
8.1	Connecting to the Network Interconnect	8-1
8.2	Accessing Important Documentation	8-2
8.3	Ensuring Viability of the Write-Back Cache	8-2
8.4	Ensuring VMScluster System Compatibility	8-2
8.5	Activating Shadowing Functionality on the Network Array	8-3
8.6	Installing and Registering Licenses	8-3

8.6.1	Setting Up the License Database on a Cluster	8-4
8.6.1.1	Common License Database	8-4
8.6.1.2	Separate License Databases	8-5
8.6.2	Registering the Licenses	8-5

9 Supporting and Operating the StorageWorks Network Storage Array

9.1	Customizing the System	9-1
9.1.1	Customizing System Files	9-1
9.1.2	Configuring and Starting DECnet for OpenVMS Alpha Software	9-2
9.1.3	Installing Layered Products	9-3
9.2	Booting a Server Processor	9-3
9.3	Backing Up the Network Array System Disk	9-4
9.4	Defining Packet Sizes	9-7
9.5	Running AUTOGEN	9-8
9.6	Balancing I/O Load on the Dual-Processor Network Arrays	9-8
9.6.1	Device-Level Balancing	9-9
9.6.2	Mount Path-Level Balancing via the DCL PREFER Command	9-9

10 System Software Maintenance

10.1	Restoring the System Disk from a Backup	10-1
10.2	Rebuilding the System Disk from the Network Array CD-ROM	10-3

A Network Storage Array NFS Matrix

B Network Storage Array Interconnect Cables

StorageWorks Glossary

Index

Examples

3-1	Server Processor Self Test Display	3-12
4-1	Software Customization Procedure Log File—Part 1	4-6
4-2	Software Customization Procedure Log File—Part 2	4-9

Figures

1-1	SWXNA-Bx/Gx Diagram	1-2
1-2	SWXNA-Bx/Gx with SWXNA-Cx Diagram	1-3
1-3	SWXNA-Ax/Fx Diagram	1-4
1-4	SWXNA-Ex/JxDiagram	1-5
1-5	StorageWorks Network Storage Array Major Components	1-10
1-6	Shadowed System Disk	1-12
1-7	StorageWorks Network Storage Array OCP	1-13
2-1	StorageWorks Network Storage Array Primary Power Plugs	2-1
2-2	Minimum Installation Clearances	2-3
3-1	Shipping Container Contents	3-3

3-2	Shipping Pallet Ramp Installation	3-5
3-3	Shipping Bolts and Brackets	3-6
3-4	Removing the Cabinet from the Pallet	3-7
3-5	Leveler Foot Adjustment	3-8
3-6	Cabinet Skirt Installation	3-9
3-7	Connecting to the Terminal Port of the StorageWorks Network Storage Array	3-11
5-1	Example Cluster Information	5-2
5-2	CFMENU Main Menu	5-3

Tables

1	StorageWorks Related Documentation	viii
2	Acoustics—Preliminary Declared Values per ISO 9296 and ISO 7779	xii
3	Schallemissionswerte—Vorläufige Werteangaben nach ISO 9296 und ISO 7779/DIN EN27779	xii
1-1	StorageWorks Network Storage Array Upgrades	1-6
1-2	StorageWorks Network Storage Array Specifications	1-7
1-3	SWXNA-Bx/Gx/SWXNA-Cx Specifications	1-8
1-4	StorageWorks Network Storage Array OCP Controls and Indicators	1-13
3-1	Thermal Stabilization Specifications	3-2
4-1	Software Configuration Tasks	4-1
4-2	Required Information for Customizing the OpenVMS Alpha Operating System	4-4
5-1	CFMENU Information Headings	5-4
7-1	Checklist for Configuring the NFS Server	7-3
B-1	StorageWorks Network Storage Array Interconnect Cables	B-1

Preface

This manual describes the procedures necessary to install, configure, and use the StorageWorks™ Network Storage Array. Mechanical cabinet configuration and power configuration as well as software installation, configuration, and upgrade procedures are covered in this manual.

Intended Audience

This manual is intended for use by customers and Digital™ Multivendor Customer Services personnel responsible for installing and configuring a StorageWorks Network Storage Array.

Structure

This manual is organized as follows:

Chapter 1	Provides an overview of the StorageWorks Network Storage Arrays and a description of their major components.
Chapter 2	Provides site preparation information.
Chapter 3	Describes the unpacking and placement of a StorageWorks Network Storage Array.
Chapter 4	Describes installing StorageWorks Network Storage Array software.
Chapter 5	Describes configuring the StorageWorks Network Storage Array storage devices.
Chapter 6	Describes the Network Storage Array's embedded read cache.
Chapter 7	Describes configuring the Network File System.
Chapter 8	Describes how to connect the network array to the network interconnect and prepare for operations.
Chapter 9	Describes supporting and operating the StorageWorks Network Storage Array in a VMSccluster™ system.
Chapter 10	Describes system software maintenance tasks.
Glossary	Defines the acronyms and specialized terms used in the StorageWorks environment.

Related Documents

Table 1 lists the StorageWorks-related user documents organized by use, system, or product. It also lists the OpenVMS Alpha™ operating system-related documents.

Table 1 StorageWorks Related Documentation

Document Title	Order Number
StorageWorks Primary Publications	
<i>StorageWorks Solutions Configuration Guide</i>	EK-BA350-CG
<i>StorageWorks Solutions Shelf and SBB User's Guide</i>	EK-BA350-UG
<i>StorageWorks Solutions BA356-SB 16-Bit Shelf User's Guide</i>	EK-BA356-UG
<i>StorageWorks Array Controllers HS Family of Array Controllers User's Guide</i>	EK-HSFAM-UG
StorageWorks Network Storage Array Publications	
<i>StorageWorks Solutions StorageWorks Network Storage Array Server Processor Component Manual</i>	EK-SWXNA-CM
StorageWorks Enclosures	
<i>StorageWorks Solutions SW500 and SW800 Cabinet Metric Shelf Bracket Kit Installation Guide</i>	EK-35XRD-IG
<i>StorageWorks RETMA Shelf Rail Kit Installation Guide</i>	EK-35XRB-IG
<i>StorageWorks Solutions SW800-Series Data Center Cabinet Cable Distribution Unit Installation Guide</i>	EK-SWCDU-IS
<i>StorageWorks Solutions SW800-Series Data Center Cabinet Installation and User's Guide</i>	EK-SW800-IG
Storage Devices	
<i>StorageWorks Building Blocks User's Guide</i>	EK-SBB35-UG
<i>StorageWorks Solutions 3½-Inch Storage Device Installation Guide</i>	EK-MC350-IG
<i>StorageWorks Solutions 5¼-Inch Storage Device Installation Guide</i>	EK-MC525-IG
DEC TCP/IP Services for OpenVMS Publications	
<i>DEC TCP/IP Services for OpenVMS Installation and Configuration</i>	AA-LU49G-TE
<i>DEC TCP/IP Services for OpenVMS Management</i>	AA-LU50G-TE
<i>DEC TCP/IP Services for OpenVMS Concepts and Planning</i>	AA-Q06TC-TE
<i>DEC TCP/IP Services for OpenVMS Management Command Reference</i>	AA-PQQGD-TE

(continued on next page)

Table 1 (Cont.) StorageWorks Related Documentation

Document Title	Order Number
OpenVMS Alpha Operating System Related Documentation	
<i>A Comparison of System Management on OpenVMS Alpha and OpenVMS VAX</i>	AA-PV71B-TE
<i>OpenVMS System Manager's Manual: Essentials</i>	AA-PV5MC-TK
<i>DECnet for OpenVMS Guide to Networking</i>	AA-PV5ZA-TK
<i>OpenVMS License Management Utility Manual</i>	AA-PVXUC-TK
<i>VMScluster Systems for OpenVMS</i>	AA-PV5WB-TK
<i>OpenVMS Alpha Version 6.2 Upgrade and Installation Manual</i>	AA-PV6XC-TE
<i>Guide to OpenVMS Performance Management</i>	AA-PV5XA-TE
<i>OpenVMS System Manager's Manual: Tuning, Monitoring and Complex Systems</i>	AA-PV5NC-TK
<i>OpenVMS Alpha Systems Dump Analyzer Utility Manual</i>	AA-PV6UB-TE
<i>Volume Shadowing for OpenVMS</i>	AA-PVXMC-TE
<i>OpenVMS System Management Utilities Reference Manual: A-L</i>	AA-PV5PC-TK
<i>OpenVMS System Management Utilities Reference Manual: M-Z</i>	AA-PV5QC-TK
General Reference Publications	
<i>Digital Systems and Options Catalog</i>	†
† Available from your Digital account representative.	

Documentation Conventions

The following conventions are used in this manual:

- | | |
|----------------------|--|
| boldface type | Boldface type indicates the first instance of terms being defined in text, in the glossary, or both. |
| <i>italic type</i> | Italic type indicates emphasis and complete manual titles. In the glossary, italic type also is used to indicate cross-references. |

Manufacturer's Declarations

Following are manufacturer's declarations applicable to the StorageWorks Network Storage Array:

CAUTION

This is a class A product. In a domestic environment, this product may cause radio interference, in which case the user may be required to take adequate measures.

ACHTUNG !

Dieses ist ein Gerät der Funkstörgrenzwertklasse A. In Wohnbereichen können bei Betrieb dieses Gerätes Rundfunkstörungen auftreten, in welchen Fällen die Benutzer für entsprechende Gegenmaßnahmen verantwortlich sind.

ATTENTION !

Ceci est un produit de Classe A. Dans un environnement domestique, ce produit risque de créer des interférences radiélectriques, il appartiendra alors à l'utilisateur de prendre les mesures spécifiques appropriées.

Note

The equipment described in this manual is listed by the Underwriters Laboratories Incorporated and bears the UL Listing mark. The StorageWorks Network Storage Array cabinets also are certified by the Canadian Standards Association and TUV Product Service GmbH and bear both the CSA certification and TUV GS marks.

Table 2 Acoustics—Preliminary Declared Values per ISO 9296 and ISO 7779

Product†	Sound Power Level L_{WAd} , B‡		Sound Pressure Level L_{pAm} , dBA (Bystander Positions)	
	Idle	Operate	Idle	Operate
SWXNA–xx with only cabinet fans operating	7.6	7.6	59	59
SWXNA–xx with 2 BA350–M shelves and 12 BA350–S shelves, each containing 6 RZ29–VA disk drives	7.6	7.6	59	59
Per device when installed in SWXNA				
BA350–S shelf containing 6 RZ29–VA disk drives	5.7	5.7	39	39
BA350–M shelf	5.6	5.6	39	39

† Current values for specific configurations are available from Digital representatives.
‡ 1 B = 10 dBA.

Note

Table 3 is a translation of the English language specifications in Table 2 into the German language.

Table 3 Schallemissionswerte—Vorläufige Werteangaben nach ISO 9296 und ISO 7779/DIN EN27779

Gerät†	Schalleistungspegel L_{WAd} , B‡		Schalldruckpegel L_{pAm} , dBA (Beistehende Position)	
	Leerlauf	Betrieb	Leerlauf	Betrieb
SWXNA–xx nur mit kabinett Lüftern in Betrieb	7,6	7,6	59	59
SWXNA–xx mit 2 BA350–M shelves und 12 BA350–S shelves, jedes bestückt mit 6 RZ29–VA disk drives	7,6	7,6	59	59
Pro Gerät installiert im SWXNA				
BA350–S shelf mit 6 RZ29–V disk drives	5,7	5,7	39	39
BA350–M shelf	5,6	5,6	39	39

† Aktuelle Werte für spezielle Ausrüstungsstufen sind über die Digital Equipment Vertretungen erhältlich.
‡ 1 B = 10 dBA.

This chapter presents a system overview of the StorageWorks Network Storage Array, describes each of its major components, and lists the upgrade options.

1.1 System Overview

The StorageWorks Network Storage Arrays integrate Digital's Alpha technology with StorageWorks modular design to meet the storage requirements of large FDDI/CDDI-based network systems. The StorageWorks Network Storage Arrays support a wide range of StorageWorks solid state disk, magnetic disk, tape, optical, and loader devices for configuring the exact storage solution to meet your application's needs. The StorageWorks Network Storage Array family extends the StorageWorks line with both standard and highly available servers to meet the I/O needs of even the largest network systems. Each StorageWorks Network Storage Array comprises part of a network environment.

The following StorageWorks Network Storage Arrays currently are available:

- SWXNA-BA, BB, GA, GB
- SWXNA-AA, AB, FA, FB
- SWXNA-EA, EB, JA, JB

The following expansion options also are available:

- SWXNA-CA, CB
- SWXNA-DA, HA

In addition, many kits to upgrade your network array are available. See the *Digital Systems and Options Catalog* for a complete list.

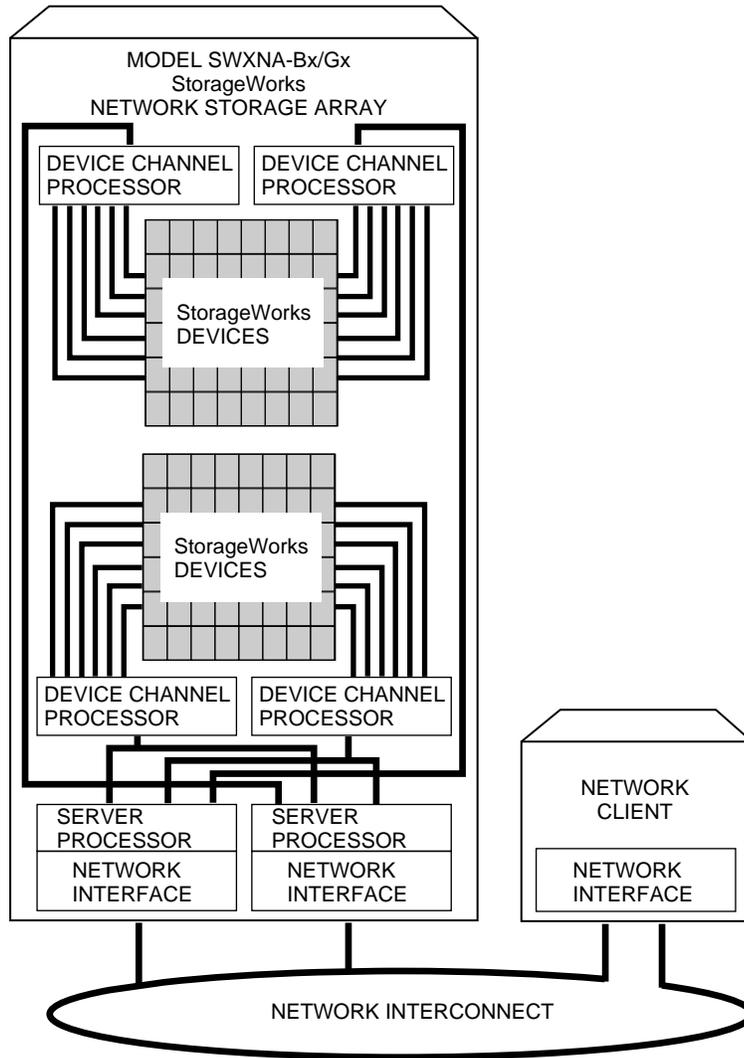
1.1.1 SWXNA-Bx/Gx Description

Figure 1-1 shows a conceptual model of the SWXNA-Bx/Gx StorageWorks Network Storage Array. Ideal for large networks running "bet-your-business" applications, the SWXNA-Bx/Gx network array provides high availability via dual-redundant server processors and four dual-redundant device channel processors. Housed in a single StorageWorks SW800 cabinet, this network array can connect up to 72 redundant devices. Maximum SWXNA-Bx/Gx online disk capacity is 309 GB using RZ29 disk drives. Consult the Software Product Description (SPD), *HS1CP Device Channel Processor Operating Firmware Version 2.x*, SPD 64.19.xx, for a complete list of supported StorageWorks devices. Nearline storage of 10-plus terabytes is possible by adding up to two StorageWorks tape libraries for backups and archiving.

During normal operation, I/O performance is enhanced by balancing the I/O load across the redundant system components in a customer-definable manner. If a component failure occurs, the redundant partner takes over, providing continued service until the failed component can be repaired or replaced. All SWXNA

redundant StorageWorks Network Storage Array components can be hot swapped or replaced without interrupting I/O service.

Figure 1-1 SWXNA-Bx/Gx Diagram

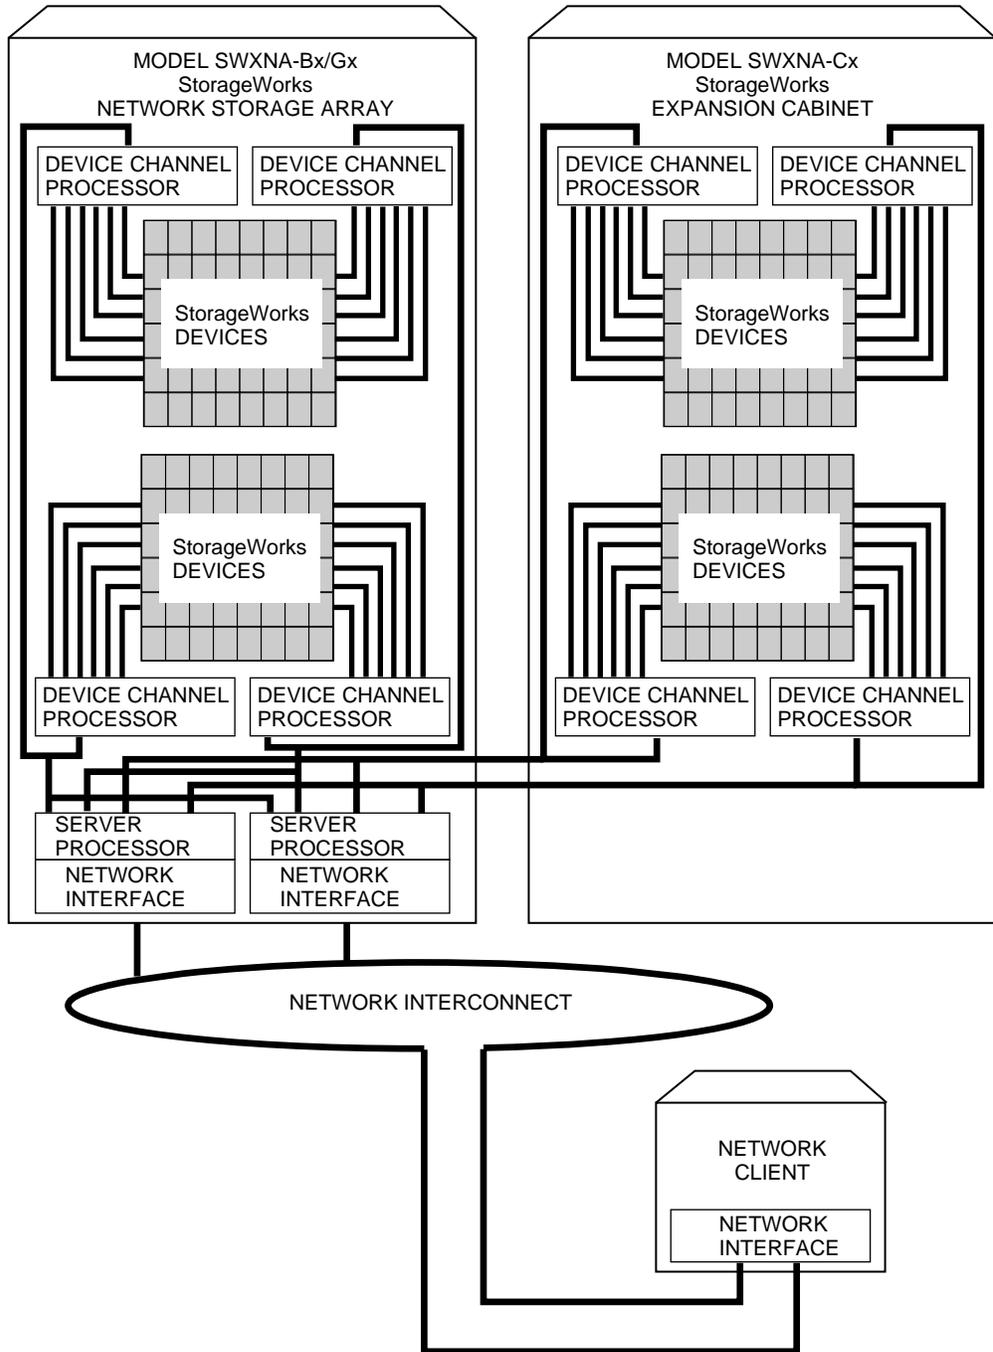


CXO-5272A-MC

1.1.2 SWXNA-Cx Description

Double the device support (to 144 devices) and online disk capacity to 619 GB by connecting the SWXNA-Bx/Gx to an optional SWXNA-Cx expansion cabinet. The SWXNA-Cx includes a second StorageWorks SW800 cabinet and four additional redundant device channel processors. Redundant power and cooling are also standard at the SWXNA-Cx cabinet level. Figure 1-2 shows a conceptual model of the SWXNA-Bx/Gx StorageWorks Network Storage Array with the SWXNA-Cx expansion option installed.

Figure 1–2 SWXNA-Bx/Gx with SWXNA-Cx Diagram



CXO-5273A-MC

1.1.3 SWXNA-Ax/Fx Description

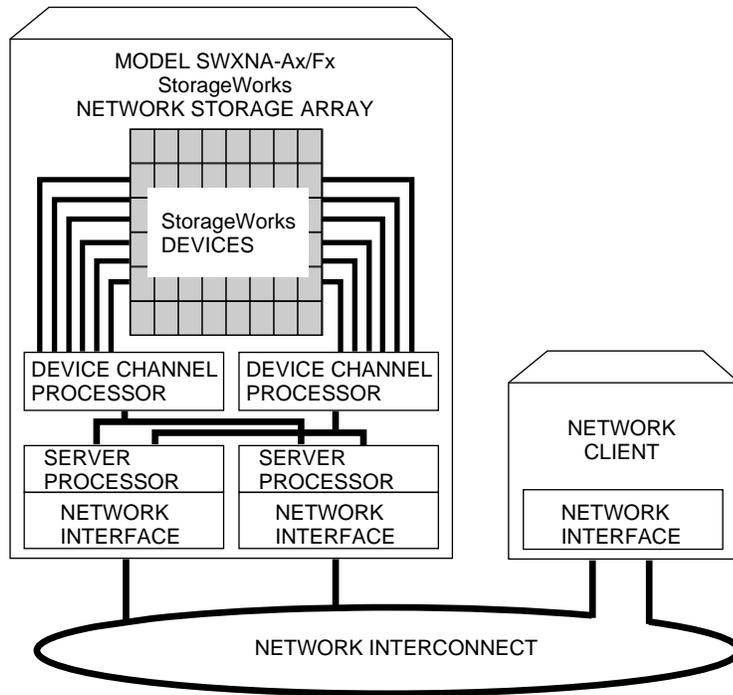
Figure 1–3 shows a conceptual model of the SWXNA-Ax/Fx StorageWorks Network Storage Array. It is a fully-redundant, high-availability storage solution offering full protection against any single component failure. These network arrays are suited to medium-sized operations running “mission-critical” applications. High network array availability is assured with fully redundant server processors, two paths to the network interface, two paths

to every storage device and standard redundant power and cooling. The full range of StorageWorks redundant power and cooling features for devices are also available.

The basic SWXNA-Ax/Fx StorageWorks Network Storage Array provides fully redundant access to up to 36 redundant StorageWorks device connections. Consult the Software Product Description, *HS1CP Device Channel Processor Operating Firmware Version 2.x*, SPD 64.19.xx, for a complete list of supported StorageWorks devices. Using the RZ29 disk drive, 154 GB of online, redundant storage per network array is available. The model SWXNA-Ax/Fx StorageWorks Network Storage Array and its attached storage devices are housed in a single StorageWorks SW800 cabinet.

During normal operation, I/O performance is enhanced by balancing the I/O load across the redundant system components in a customer-definable manner. If a component failure occurs, the redundant partner takes over, providing continued service until the failed component can be repaired or replaced. All SWXNA redundant StorageWorks Network Storage Array components can be hot swapped or replaced without interrupting I/O service.

Figure 1-3 SWXNA-Ax/Fx Diagram



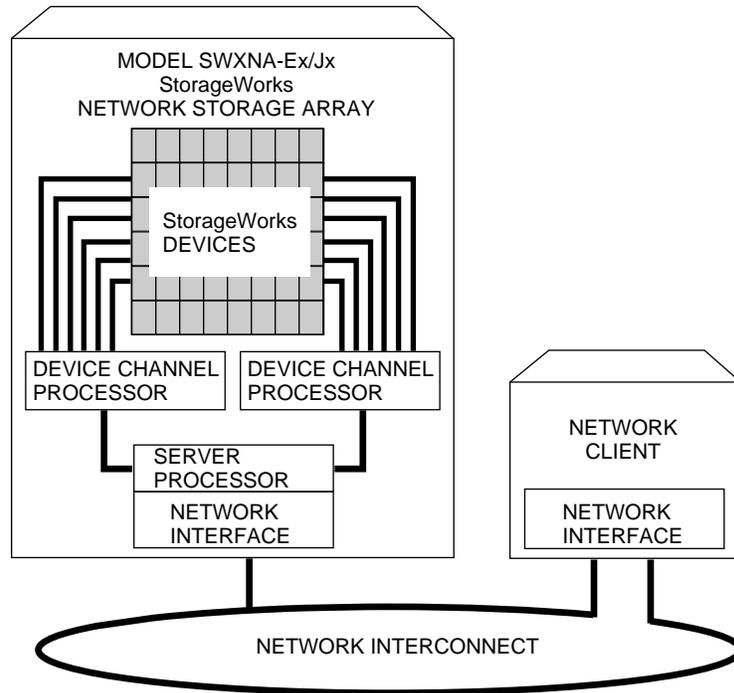
CXO-5274A-MC

1.1.4 SWXNA-Ex/Jx Description

Figure 1-4 shows a conceptual model of the SWXNA-Ex/Jx StorageWorks Network Storage Array. Suited to medium-sized operations where you need distributed access and storage management benefits, these network arrays consist of a single Alpha-based server processor and two device channel processors (HS1CPs) capable of connecting the server processor to up to 36 redundant StorageWorks devices. Consult the Software Product Description, *HS1CP Device*

Channel Processor Operating Firmware Version 2.x, SPD 64.19.xx, for a complete list of supported StorageWorks devices. Using the RZ29 disk drive, up to 154 GB of redundant online storage per network array is available. The model SWXNA-Ex/Jx StorageWorks Network Storage Array and its attached storage devices are housed in a single StorageWorks SW800 cabinet.

Figure 1–4 SWXNA-Ex/JxDiagram



CXO-5449A-MC

1.1.5 SWXNA-DA/HA Description

The SWXNA-DA/HA option kit contains all the major components of an SWXNA-Ex/Jx except for the StorageWorks SW800 cabinet. These options provide for the conversion of an existing SW800 cabinet to a StorageWorks SWXNA-Ex/Jx Network Storage Array, and are described in the *SWXNA-DA/HA Installation Guide*.

1.1.6 SWXNA-xx Series Software and Firmware Components

The SWXNA-xx series server processor runs the OpenVMS Alpha operating system. This means that, in addition to serving online storage to hosts, the StorageWorks Network Storage Array can use a wide range of StorageWorks and POLYCENTER™ storage management software to provide an optimally managed storage environment.

StorageWorks Network Storage Arrays support block I/O services and, used with TCP/IP services for OpenVMS Alpha software, support NFS (Network File Server) services. A distributed, mixed environment that can include OpenVMS clusters, CI clusters, UNIX® workstations, and PCs can be supported for any system with NFS Version 2 client capabilities.

As full members of a network, StorageWorks Network Storage Arrays merge seamlessly into existing clusters and can evolve in parallel with or independent of the operating system. The StorageWorks Network Storage Arrays can manage storage services to as many as 95 other host computers or workstations in a network system.

To further enhance both performance and data availability, host-based RAID or the optional controller-based RAID can be added to the network array in addition to the standard, nonvolatile write-back cache. Consult the Software Product Description, *HS1CP Device Channel Processor Optional Operating Firmware, Version 2.x*, SPD 64.20.xx, for further information.

1.2 Available Upgrades

The StorageWorks Network Storage Arrays have a number of applicable upgrades: some are StorageWorks Network Storage Array specific, others are applicable to the SW800 enclosure or StorageWorks shelves and support the StorageWorks Network Storage Array upgrades. Each upgrade is briefly described in the Table 1–1.

Table 1–1 StorageWorks Network Storage Array Upgrades

Upgrade Order Number	Description
StorageWorks Network Storage Array-Specific Upgrades	
SWXNA–CA/CB	StorageWorks Network Storage Array expansion cabinet for attachment to SWXNA-Bx/Gx. Expands the number of device buses to 8 and available device ports to 144. Available in 50 and 60 Hz options.
SWXNA–DA/HA	Expands an existing SW800 cabinet to an SWXNA-Ex/Jx, providing nonredundant access to the installed StorageWorks devices. Available in FDDI or CDDI options.
SWXUA–E1	Upgrades an HS111 to an SWXNA-EA/EB.
SWXUA–A1	Upgrades an HS121 to an SWXNA-Ax; available in 50 and 60 Hz options.
SWXUA–B1	Upgrades an HS121 to an SWXNA-Bx; available in 50 and 60 Hz options.
SWXUA–E2	Upgrades an HS211 to an SWXNA-EA/EB; available in 50 and 60 Hz options.
SWXUA–A2	Upgrades an HS221 to an SWXNA-Ax. Available in 50 and 60 Hz options.
SWXUA–B2	Upgrades an HS241 to an SWXNA-Bx. Available in 50 and 60 Hz options.
SWXUA–AA	Upgrades an SWXNA-EA/JA to an SWXNA-AA/FA. Available for both FDDI and CDDI.
SWXUA–AB	Upgrades an SWXNA-EB/JB to an SWXNA-AB/FB. Available for both FDDI and CDDI.
SWXUA–BA	Upgrades an SWXNA-Ax/Fx to an SWXNA-Bx/Gx. Available in 50 and 60 Hz options for both FDDI and CDDI.
SWXUA–M1	StorageWorks Network Storage Array 64 MB server memory expansion.

(continued on next page)

Table 1–1 (Cont.) StorageWorks Network Storage Array Upgrades

Upgrade Order Number	Description
StorageWorks Network Storage Array-Specific Upgrades	
SWXUA-M2	StorageWorks Network Storage Array 128 MB server memory expansion.
SWXUA-PA	StorageWorks Network Storage Array server processor redundant power supply.
Supporting Upgrades	
BA350-JA	StorageWorks device shelf with mounting kit and power supply. Provide mounting for up to seven StorageWorks devices.
BA35X-HA	StorageWorks device shelf power supply.
SW8XP-AX	Additional power controller for the SW800 cabinet.

1.2.1 StorageWorks Network Storage Array Specifications

Tables 1–2 and 1–3 provide the StorageWorks Network Storage Array specifications.

Table 1–2 StorageWorks Network Storage Array Specifications

	Models SWXNA-Ex/Jx	Model SWXNA-Ax/Fx
Storage Capacity		
Maximum number of disks	36 redundant connections	36 redundant connections
Maximum disk capacity using RZ29 disks	154 GB	154 GB
Maximum number of tape drives	See the <i>StorageWorks Solutions Configuration Planning Guide</i> for configuration guidelines.	
Maximum number of Solid State disk drives	See the <i>StorageWorks Solutions Configuration Planning Guide</i> for configuration guidelines.	
Power Requirements		
Voltage	100 to 240 V ac	100 to 240 V ac
Frequency	50/60 Hz	50/60 Hz
Device channel processor battery backup	Standard	Standard
Redundant Power	Optional	Standard

(continued on next page)

Table 1–2 (Cont.) StorageWorks Network Storage Array Specifications

	Models SWXNA-Ex/Jx	Model SWXNA-Ax/Fx
Environmental		
Operating temperature	10 C to 40 C	10 C to 40 C
Nonoperating temperature	-40 C to 66 C	-40 C to 66 C
Relative humidity	10% to 80%	10% to 80%
Altitude	Sea level to 2400 m (8000 ft)	Sea level to 2400 m (8000 ft)
Physical Dimensions		
Height	1700 mm (67 in)	1700 mm (67 in)
Width	800 mm (31 in)	800 mm (31 in)
Depth	875 mm (34.5 in)	875 mm (34.5 in)
Weight (without devices)	310 kg (685 lbs)	355 kg (780 lbs)

Table 1–3 SWXNA–Bx/Gx/SWXNA–Cx Specifications

	Model SWXNA–Bx/Gx	Model SWXNA–Bx/Gx with SWXNA–Cx Installed
Storage Capacity		
Maximum number of disks	72 redundant connections	144 redundant connections
Maximum disk capacity using RZ29 disks	309 GB	619 GB
Maximum number of tapes	See the <i>StorageWorks Solutions Configuration Guide</i> for configuration guidelines.	
Maximum number of Solid State disks	See the <i>StorageWorks Solutions Configuration Guide</i> for configuration guidelines.	
I/O Performance		
Saturation performance: single-block disk I/O request rate	4,300	4,300
Saturation performance: Maximum data transfer rate (to/from network interconnect)	11 MB/second	11 MB/second
I/O response time for typical OpenVMS load	12 ms (cache disabled) 7 ms (read and write cache enabled)	12 ms (cache disabled) 7 ms (read and write cache enabled)

(continued on next page)

Table 1–3 (Cont.) SWXNA–Bx/Gx/SWXNA–Cx Specifications

	Model SWXNA–Bx/Gx	Model SWXNA–Bx /Gx with SWXNA–Cx Installed
Power Requirements		
Voltage	100 to 240 V ac	100 to 240 V ac
Frequency	50/60 Hz	50/60 Hz
Device channel processor battery backup	Standard	Standard
Redundant Power	Optional	Standard
Environmental		
Operating temperature	10 C to 40 C	10 C to 40 C
Nonoperating temperature	-40 C to 66 C	-40 C to 66 C
Relative humidity	10% to 80%	10% to 80%
Altitude	Sea level to 2400 m (8000 ft)	Sea level to 2400 m (8000 ft)
Physical Dimensions		
Height	1700 mm (67 in)	1700 mm (67 in)
Width	800 mm (31 in)	800 mm (31 in)
Depth	875 mm (34.5 in)	875 mm (34.5 in)
Weight (without devices)	310 kg (685 lbs)	355 kg (780 lbs)

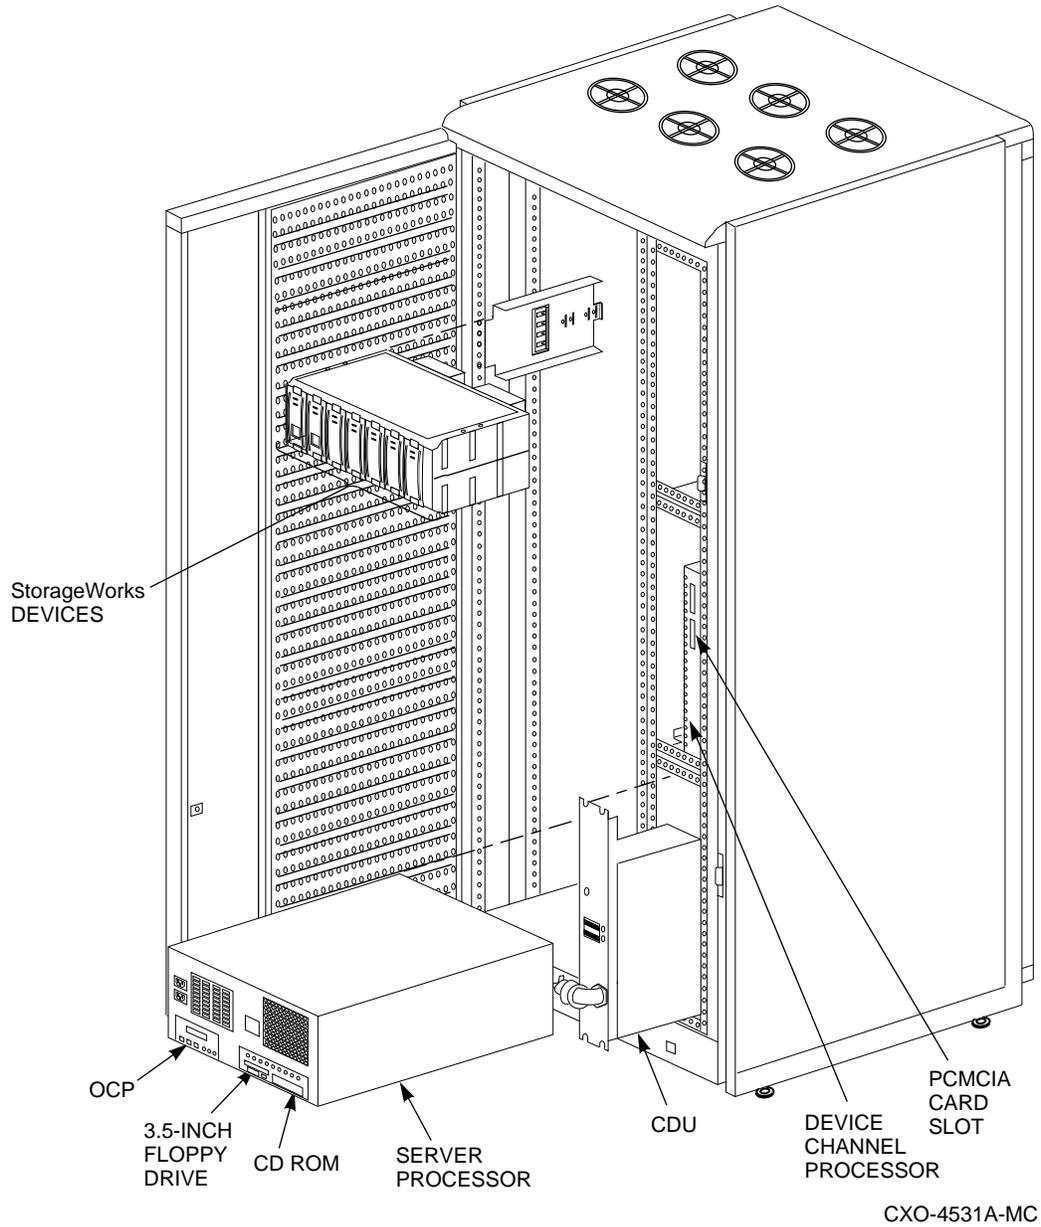
1.3 Major Components

Figure 1–5 shows a StorageWorks Network Storage Array calling out each of the major components and subcomponents. The major components are as follows:

- Server processor
- Device channel processor (HS1CP)
- StorageWorks storage devices

The subcomponents are explained in Section 1.3.1 through Section 1.3.3.

Figure 1-5 StorageWorks Network Storage Array Major Components



1.3.1 Server Processor

The StorageWorks Network Storage Array's server processor combines Digital's advanced Alpha technology with 128 MB RAM to provide the internal performance required to handle the I/O load of network systems.

Each server processor supports the full range of OpenVMS standard and optional storage management utilities, providing a managed storage environment that the OpenVMS system administrator will find familiar.

In the model SWXNA-Ax/Bx/Fx/Gx StorageWorks Network Storage Array, two server processors share the I/O load under normal circumstances. When a failure occurs, the operational server processor takes over the load of its failed partner and continues to deliver data from all devices until repair is effected. Because server processors are separately powered and cooled, a power supply or cooling failure does not interrupt data access.

1.3.1.1 Shadowed System Disk

SCSI hard drives configured in a shadowset to provide redundancy are used as the system disk for the server processors. The system disk contains the OpenVMS Alpha operating system, the Software Customization Procedure (SCP), various relevant documents, and required layered software components. Figure 1-6 shows the location of the shadowed system disk and the cabling to the PCI to SCSI adapters in the server processors.

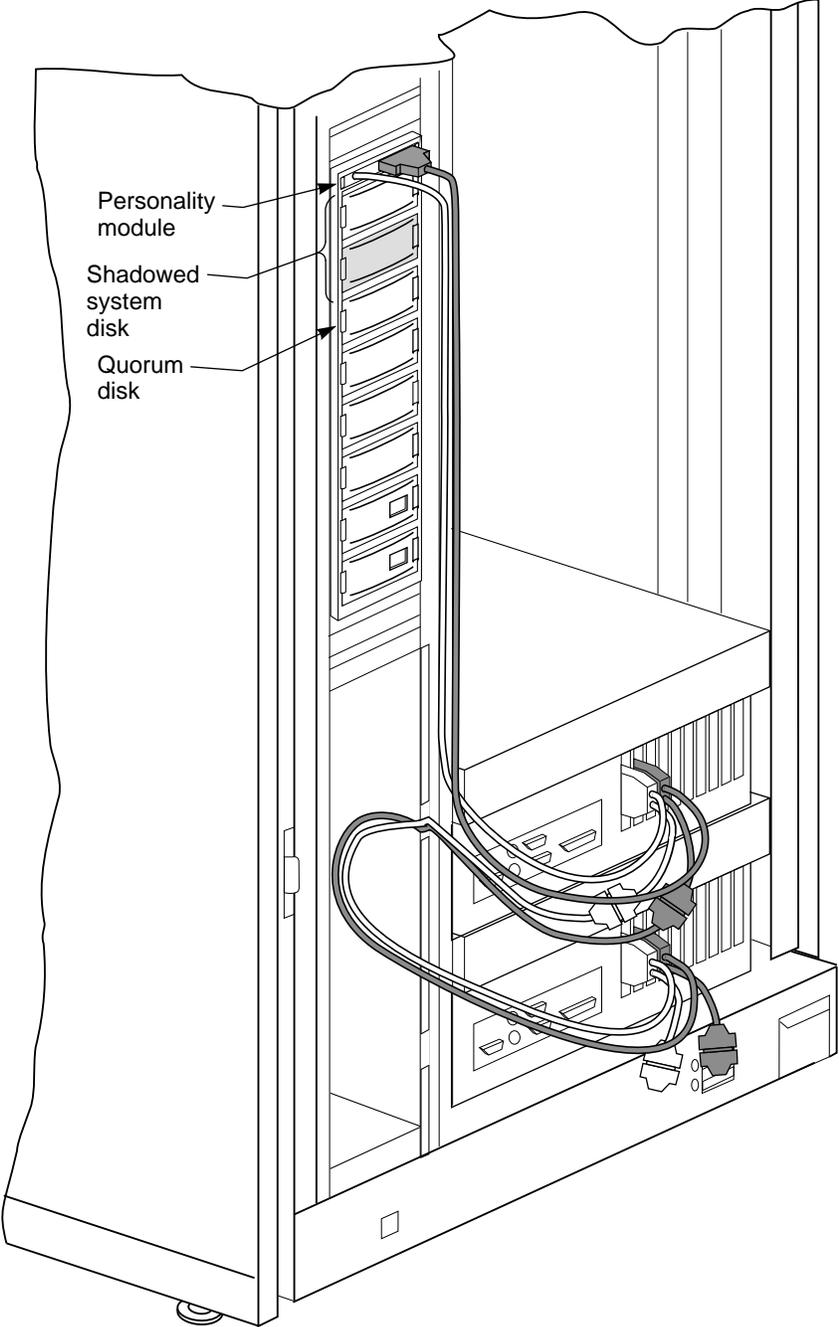
1.3.1.2 CD-ROM Drive

The CD-ROM drive is used to mount the two CD-ROMs and to access files for various purposes. The StorageWorks Network Storage Array includes two CD-ROMs, an operating system CD-ROM and a documentation CD-ROM.

The operating system CD-ROM is used for system rebuild operations if there is no system disk backup, for updating the network array operating system to a new version of the OpenVMS Alpha operating system, or for operations performed directly from the CD-ROM such as backup. The documentation CD-ROM is used to access documentation related to the OpenVMS Alpha operating system.

For more detailed information on the use of these two CD-ROMs, see Section 1.4.

Figure 1-6 Shadowed System Disk



CXO-5285A-MC

1.3.1.3 3.5-Inch Floppy Drive

The 3.5-inch floppy drive is used to run the EISA Configuration Utility (ECU) whenever configuration changes are made to the server processor's internal bus, such as during a hardware upgrade.

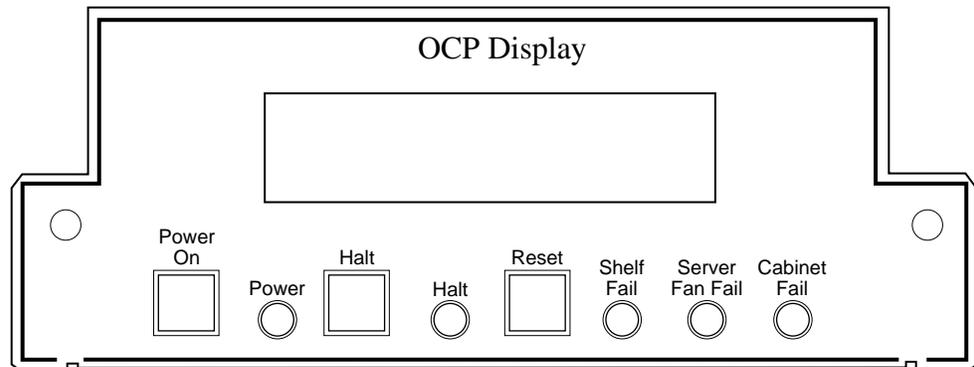
1.3.1.4 Operator Control Panel

The operator control panel (OCP) gives you another interface to the StorageWorks Network Storage Array. It provides information about and allows you to control the operation of the server processor.

Description of OCP Controls and Indicators

Figure 1-7 shows the OCP. Table 1-4 describes the controls and indicators on the OCP.

Figure 1-7 StorageWorks Network Storage Array OCP



CXO-4586A-MC

Table 1-4 StorageWorks Network Storage Array OCP Controls and Indicators

Control or Indicator	Description
OCP Display	This display provides server processor status information.
Power On switch	This switch applies power to and removes power from the server processor.
Power indicator	This indicator shows that power is enabled to the server processor when lighted.
Halt switch	This switch halts the server processor.
Halt indicator	This indicator shows that the server processor has been halted.
Reset switch	This momentary switch resets the server processor.
Shelf Fail indicator	This indicator is not used.
Server Fan Fail indicator	This indicator shows failure of fans 1 or 2 in the server processor when lighted. See Server Fan Fail Indicator for more information.
Cabinet Fail indicator	This indicator is not used.

Cycling Power with the OCP

During the installation procedure, you are directed to cycle (remove, then restore) power to the server processor. Use the Power On switch on the OCP to perform this step when directed to cycle the server processor.

Server Fan Fail Indicator

The Server Fan Fail indicator lights when either fan 1 or fan 2 fails. Periodically check to make sure the Server Fan Fail indicator is not lighted.

Failure of fan 3 (the fan cooling the microprocessor) causes the server processor to shut down. The indicator will not light in this instance. In the event of a fan 3 failure, the following message appears on the server processor terminal:

```
System Machine Check Through Vector 00000660
logout frame address 0x6048 code 0x223
Machine Check Code--> 0x223 Fan 1 Failure
```

Note

Fan 3 is the hardware designation for the fan which provides the primary cooling air to the microprocessor. This fan is designated as fan 1 to the server processor software. Because of this, the error message shown indicates fan 1 failure.

In the event of a fan failure, call your service representative.

1.3.2 Device Channel Processor

The device channel processor can provide nonredundant connections between a server processor and up to 42 industry-standard SCSI-2 devices on six independent fast (10 MB/second) SCSI-2 buses. The device channel processor supports a wide range of StorageWorks magnetic, solid state, and optical disks, tapes, and media loader devices.

To support the I/O demands on this much storage, each device channel processor is capable of processing up to 1000 I/O requests per second. The device channel processor firmware automatically balances striped I/O across member disks as well as enabling the following features:

- Redundant access
- Read/Write caching
- RAIDset configuration support

Each of these features is described later in this section.

The device channel processor firmware resides on a PCMCIA card that ships with the StorageWorks Network Storage Array. The PCMCIA card is located in the slot located on the front of the device channel processor.

1.3.2.1 Read/Write Cache

Each device channel processor includes a nonvolatile read/write cache that reduces I/O request execution times. Firmware allocates the cache dynamically based on actual I/O activity. For optimal resource utilization, management facilities allow you to enable or disable both read and write caching for each storage unit attached to a device channel processor. In the event of a power failure, on-board batteries keep write cache contents intact for up to 100 hours, so no data is lost.

Note

The write-back cache module installed in your StorageWorks Network Storage Array contains batteries that were completely charged at the factory. It is normal for these batteries to discharge slightly in shipment. The network array's write-back cache and RAID features require fully-charged batteries to maintain absolute data integrity. After installation, these advanced features may not be available until the batteries have had an opportunity to completely recharge. The charging process may take up to 4 hours to complete.

1.3.2.2 RAIDset Configuration Support

The device channel processor provides for Digital's implementation of parity RAID technology, which dynamically adjusts to the I/O workload, optimizing I/O performance in both high data-transfer rate and high transaction-rate environments. Using distributed parity, the StorageWorks Network Storage Array RAID option economically provides continued access to data in the event of a disk failure. In the fully redundant model StorageWorks Network Storage Array, this increases data availability to a level comparable to that of shadowing at significantly lower cost.

1.3.3 StorageWorks Storage Devices

Because StorageWorks Network Storage Array support a wide selection of Digital's StorageWorks components, storage subsystems can be custom configured to meet virtually any application need. For increased reliability, the SW800 cabinet that houses both StorageWorks Network Storage Array models supports the full suite of StorageWorks redundant power and cooling features.

StorageWorks enclosure components and supported storage devices are all available on a configure-to-order basis. Consult the Software Product Description for the HS1CP device channel processor for an up-to-date list of supported StorageWorks devices.

1.4 Using the CD-ROM

The StorageWorks Network Storage Array ships with the current version of the OpenVMS Alpha operating system installed on the system disk; consequently, the operating system CD-ROM is not needed for installation or for day-to-day operation. However, it is needed to perform upgrades of the system disk or to restore or rebuild the system disk in the event of a system disk failure if no backup is available. It is also useful in performing complete backups of the network array's system disk. A simple standalone-mode menu, executed from the operating system CD-ROM, offers you a quick and easy way to perform the appropriate task.

The operating system CD-ROM contains the following:

- Files to install the OpenVMS Alpha operating system
- A bootable subset of the OpenVMS Alpha operating system that permits a limited system boot directly from the CD-ROM
- The Software Customization Procedure (SCP)
- Applicable Software Product Descriptions (SPDs)

The documentation CD-ROM contains OpenVMS Alpha documentation in Bookreader™, PostScript®, and ASCII text formats.

1.4.1 Handling the CD-ROM

Each CD-ROM is enclosed in an antistatic envelope. Be careful to handle the disk only by the edges; fingerprints or smudges should be cleaned using a soft cloth.

To insert the CD-ROM into the drive, press the button on the right side of the drive. The CD tray should then extend out; it may require pulling to its fully extended position. Place the CD-ROM in the tray with the printed side facing up; no disc caddy is required. With the CD-ROM seated properly in the tray, press the drive button again; the tray will retract back into the drive, and the drive light will momentarily illuminate as the drive establishes the presence of the disk.

To remove the CD-ROM, press the drive button again, and replace the CD-ROM in its protective envelope.

1.4.2 Accessing Documentation

The network array includes a documentation CD-ROM containing documentation related to the OpenVMS Alpha operating system. Documents are available in the following formats:

- Files with the filetype .PS are PostScript files that are printed on a PostScript printer.
- Files with the filetype .TXT are ASCII text files that are printed or viewed on a character cell terminal.
- Files with the filetype .DECW\$BOOK are Bookreader files. You must have DECwindows Motif® software installed to read these files with Bookreader. For more information on Bookreader, refer to the *DECwindows Motif for OpenVMS Applications Guide*.

Note

Prior to accessing documentation on the CD-ROM device, you must load the documentation CD-ROM into the CD-ROM drive and mount it using a command similar to the following:

```
$ MOUNT/CLUSTER $13$DKA500:DOCVMS062
```

Some frequently used documentation for OpenVMS Alpha software is located on the StorageWorks Network Storage Array's system disk in the [DOCUMENTATION.V0xx] directory where xx represents the OpenVMS operating system software version number. For example, for Version 6.2, the directory is [DOCUMENTATION.V062]. Documentation on the operating system CD-ROM is in directory [ALPHA062.DOCUMENTATION].

Preparing Your Site

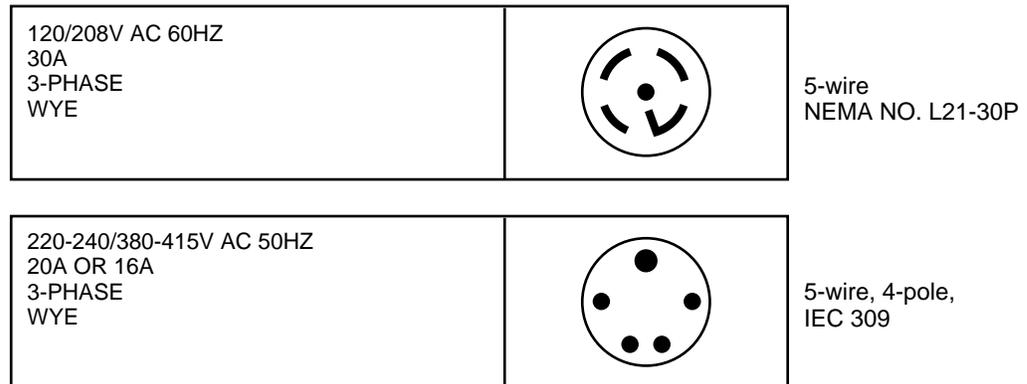
You must ensure that the site selected for installing a StorageWorks Network Storage Array satisfies requirements for power, grounding, safety, and service. This chapter discusses preparing your site to meet those requirements.

2.1 Power Requirements

The StorageWorks Network Storage Array is intended for installation in Class A computer room environments. Before installing the StorageWorks Network Storage Array, make sure that your site meets the following power requirements:

- The primary power source can supply the required amount of ac power, as specified in the specifications table (see Chapter 1).
- The proper power receptacles are installed for the cabinet. Figure 2–1 specifies the primary power plugs supplied with both 60 Hz and 50 Hz power versions of the StorageWorks Network Storage Array.

Figure 2–1 StorageWorks Network Storage Array Primary Power Plugs



CXO-3807B-MC

2.2 Grounding Requirements

The StorageWorks Network Storage Array usually is connected to other equipment by one or more network interconnect cables. For both safe and reliable operation, the cabinet and other equipment require proper grounding.

WARNING

If cabinets are not connected to a common ground, there is a potential for personal injury as a result of electric shock.

Make sure that site power distribution systems meet local electrical codes prior to the installation of StorageWorks Network Storage Array.

To make sure that the power distribution system will perform satisfactorily, a power system survey should be done before installation. The following areas should be investigated:

- Do all outlets have power ground connections?
- Do the power cords on all equipment at the site have grounding prongs?
- Are all power outlet neutral connections isolated from ground?
- Are the grounds for all outlets connected to the same power distribution panel?
- Are all devices that are connected to the same breaker as the StorageWorks Network Storage Array UL or IEC approved?

CAUTION

If there is a deficiency found in any area during the power survey, a qualified electrician must correct it before installation may begin. Failure to resolve power survey deficiencies before installing the equipment may result in personal injury as a result of electric shock.

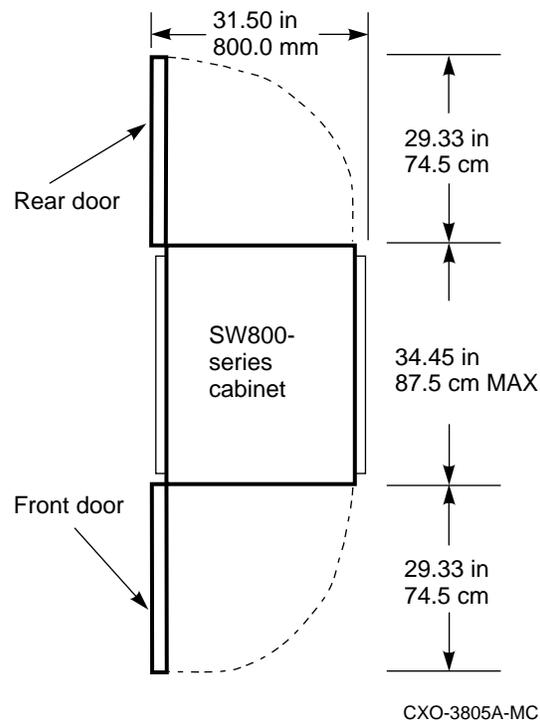
If no problems are found during the survey, the site grounding system may be considered to be adequate for safe and reliable StorageWorks Network Storage Array operation.

2.3 Safety and Service Requirements

The following general safety and service requirements refer to the floor location chosen for placing the StorageWorks Network Storage Array:

- The site floor must safely bear the weight of the cabinet, as specified in Chapter 1. The entire weight of the cabinet is supported by the four leveler feet when the cabinet is installed in its final position.
- The space around the cabinet must allow for opening the front and rear doors, for accessing cables, and for adequate airflow. See Figure 2–2 for specific space requirements.

Figure 2–2 Minimum Installation Clearances



- If the cabinet is positioned next to other cabinets, there must be sufficient service loop in any connecting cables to allow the cabinet to be moved out for access.

Note

The StorageWorks Network Storage Array is not designed to be fastened to adjacent cabinets.

Unpacking, Placing, and Applying Power to the StorageWorks Network Storage Array

This chapter describes unpacking and placing the StorageWorks Network Storage Array followed by descriptions of connecting a terminal to and applying power to the StorageWorks Network Storage Array.

3.1 Unpacking the Cabinet

The StorageWorks Network Storage Array arrives packed in a corrugated carton attached to a wooden shipping pallet, as shown in Figure 3–1. Unpack the cabinet as follows:

Note

Before unpacking the equipment, inspect the shipping carton for signs of external damage. Report any damage to the local carrier and to Digital Multivendor Customer Services or your local Digital sales office.

1. Remove the cover, the fasteners, and the corrugated board from the pallet.
2. Remove the cartons containing the ramp set and skirt kit and set them aside.
3. Cut the shipping straps. Some cabinets are packaged in a plastic or barrier bag. If the cabinet arrives in a plastic bag, leave the bag in place until the cabinet has adjusted to the local temperature and humidity, as follows:

CAUTION

Failure to thermally stabilize storage subsystems may damage drive media or associated electronics when the unit is turned on. Environmental stabilization begins when the equipment is placed in the room in which it is to be installed.

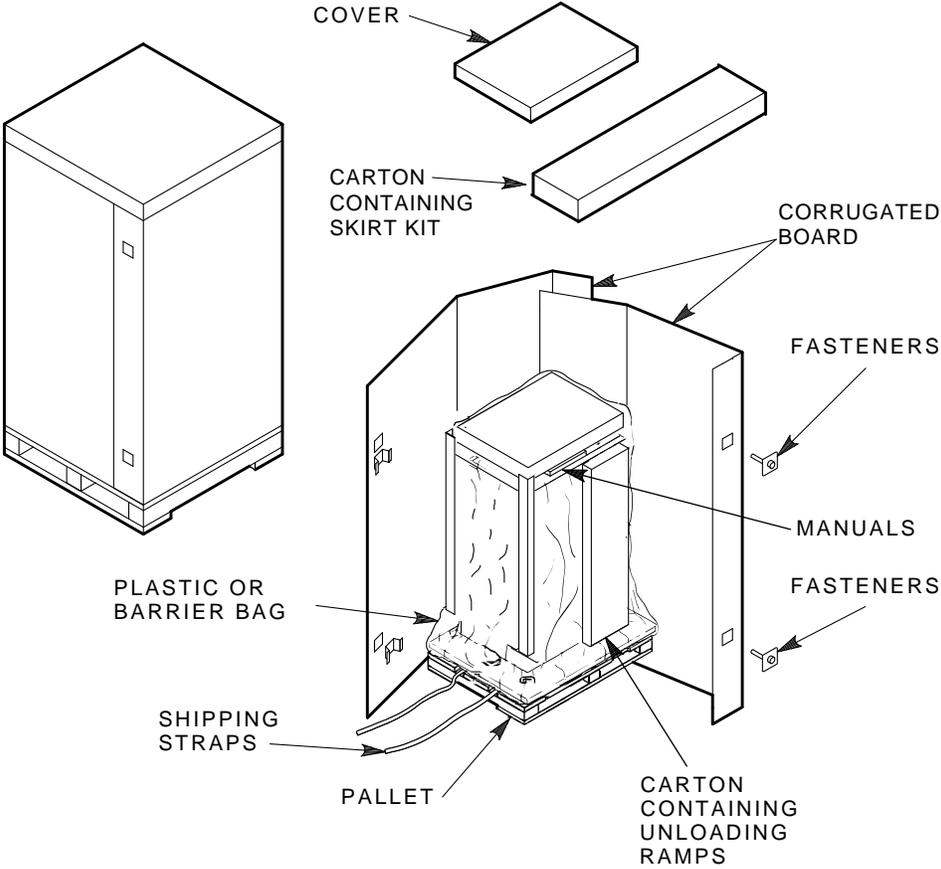
- To ensure proper operation of Digital storage devices, the StorageWorks building block (SBB) temperature must be within 18°C to 29°C (65°F to 85°F). Table 3–1 specifies the time required to thermally stabilize SBBs based on the ambient shipping temperature.

Table 3–1 Thermal Stabilization Specifications

Ambient Temperature Range °C	Ambient Temperature Range °F	Minimum Stabilization Time
60 to 66	140 to 151	3 hours
50 to 59	122 to 139	2 hours
40 to 49	104 to 121	1 hour
30 to 39	86 to 103	30 minutes
18 to 29	65 to 85	None
10 to 17	50 to 64	30 minutes
0 to 9	32 to 49	1 hour
-10 to -1	14 to 31	2 hours
-20 to -11	-4 to 13	3 hours
-30 to -21	-22 to -5	4 hours
-40 to -31	-40 to -21	5 hours

- If condensation *is visible* on the outside of the storage device, stabilize the device and the SBB in the operating environment for 6 hours or until the condensation is no longer visible, whichever is longer.
 - If condensation *is not visible* on the outside of the storage device, thermally stabilize the device for the amount of time specified in Table 3–1.
4. Once the cabinet is unpacked, examine the front and rear doors, right and left side panels, top panel, and undercarriage for any apparent damage. Report such problems immediately to Digital Multivendor Customer Services or your local Digital sales office.
 5. Retain the shipping container and all packing materials.

Figure 3-1 Shipping Container Contents



CXO-3533A-TI

3.2 Removing the Cabinet from the Pallet

WARNING

Three people are required to remove the cabinet from the shipping pallet. Failure to use sufficient personnel and correct safety precautions can result in injury and equipment damage. All personnel should wear safety glasses.

Use the following procedure to remove the cabinet from the shipping pallet:

1. Remove any packing material remaining on the pallet.
2. Remove and inspect the two unloading ramps. The ramps, ramp side rails, and metal hardware should be inspected for the following defects:
 - Cracks more than 25 percent of the ramp depth, either across or lengthwise on the ramp
 - Knots or knotholes going through the thickness of the ramp and greater than 50 percent of the ramp width
 - Loose, missing, or broken ramp side rails
 - Loose, missing, or bent metal hardware

If any of these defects exist, do not use the ramp. Investigate alternate means of removing the cabinet or order a new ramp. (The part number for the ramp set is 99-08897-05.)

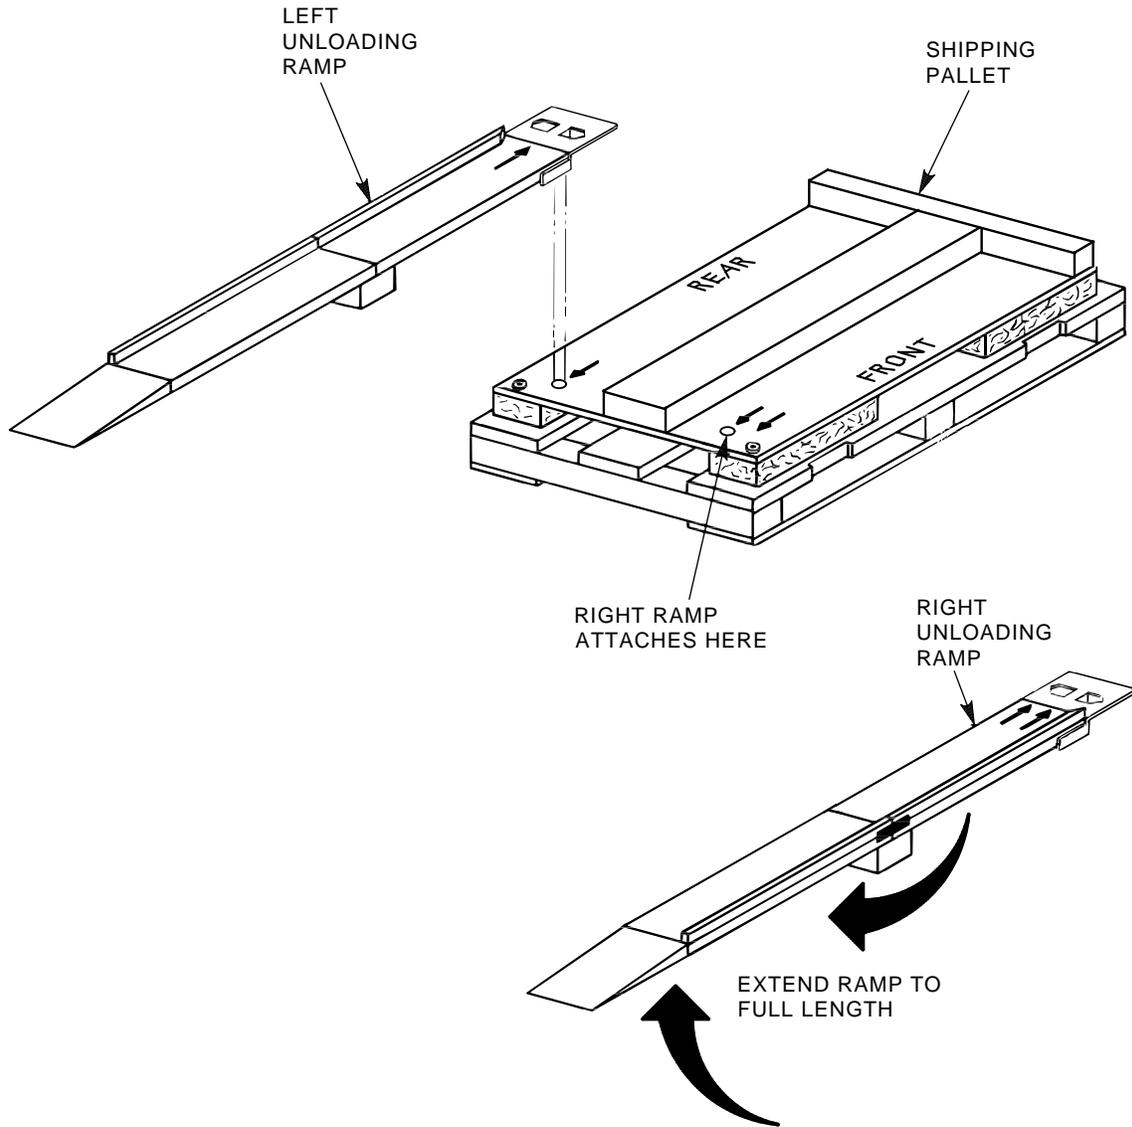
3. Attach the ramps by fitting the metal prongs into the holes on the pallet, as shown in Figure 3-2. Make sure that the arrows on the ramps match up with the arrows on the pallet.
4. Extend the ramps to their full length.
5. See Figure 3-3 for the location of the shipping bolts. Remove the bolts.
6. Remove the shipping brackets, shown in Figure 3-3, from the cabinet leveler feet and set aside.
7. Loosen the leveler locking nuts and screw the four cabinet leveler feet all the way up into the cabinet.

WARNING

The leveler feet must be raised fully for the cabinet to roll easily down the unloading ramps. Failure to do so can result in the cabinet tipping off the pallet or ramp.

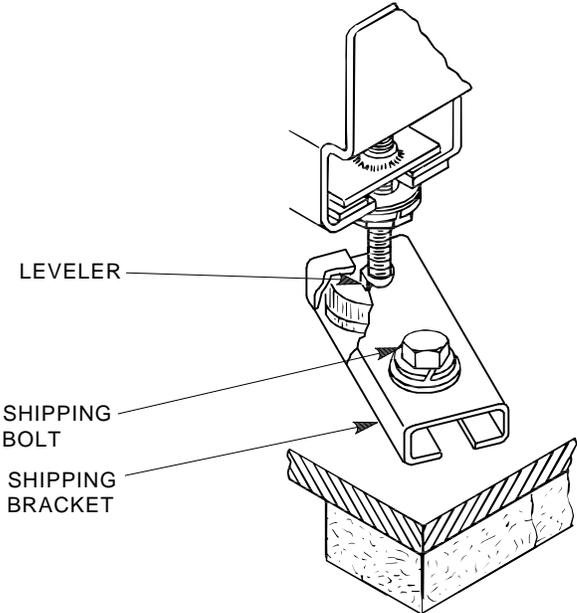
8. Carefully roll the cabinet off the pallet and down the ramps to the floor as shown in Figure 3-4.

Figure 3-2 Shipping Pallet Ramp Installation



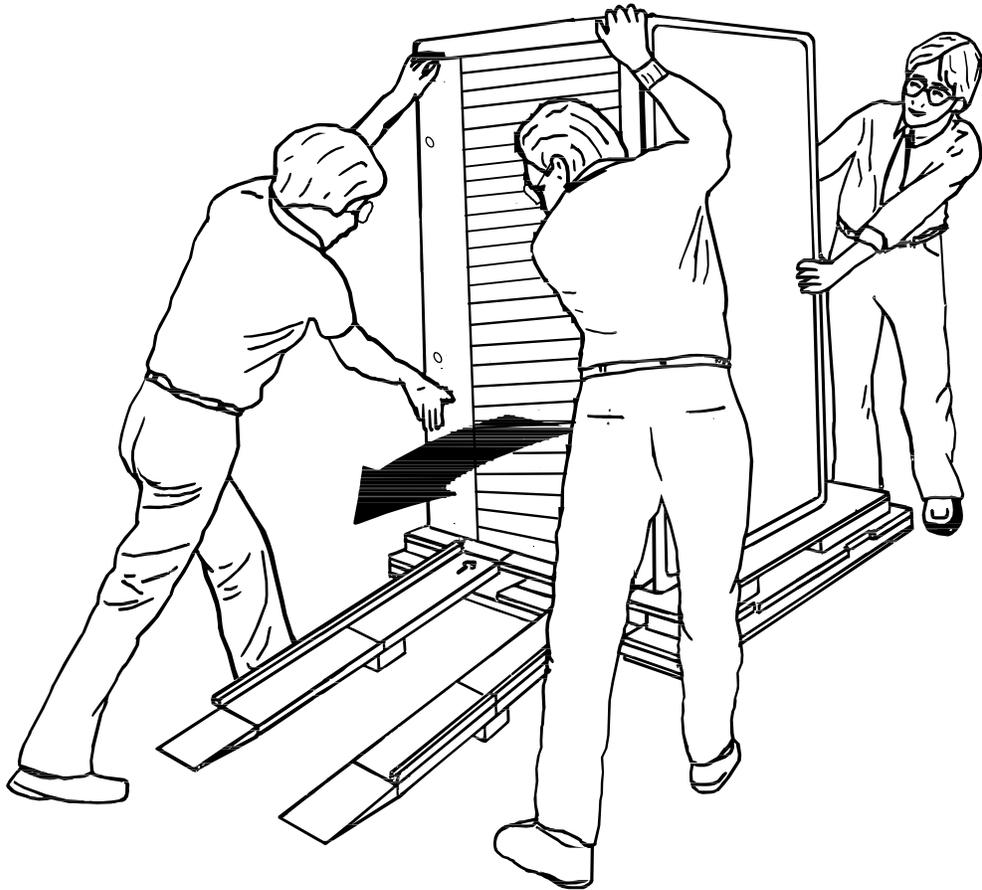
CXO-688D_S

Figure 3-3 Shipping Bolts and Brackets



SHR_X1102A_89_SCN

Figure 3–4 Removing the Cabinet from the Pallet



CXO-3808A

3.3 Placing the Cabinet

WARNING

Use extreme caution when rolling the cabinet across the floor. Failure to raise all leveler feet and to provide a clear path for the cabinet's casters can result in the cabinet tipping over and injury to personnel.

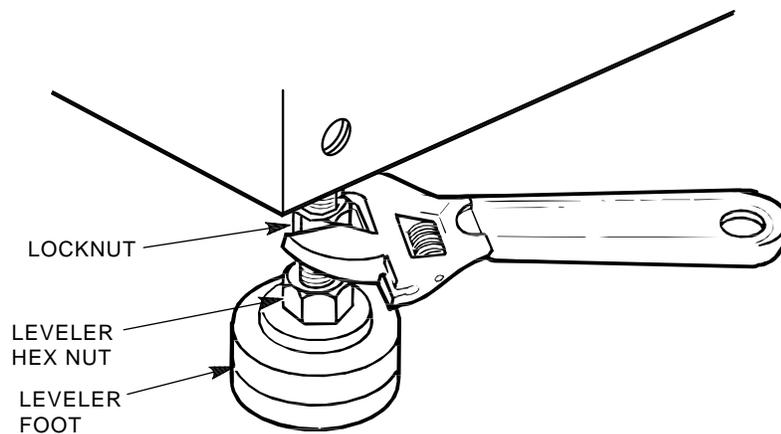
Once the cabinet has been removed from the pallet, the cabinet can be rolled to its final installation position. Secure loose cabinet cables up and out of the way when rolling the cabinet.

3.4 Leveling the Cabinet

Level the cabinet in its final position as follows:

1. Loosen the locknuts on all four leveler feet as shown in Figure 3-5.
2. Turn each leveler hex nut clockwise until the leveler foot contacts the floor.
3. Adjust all four leveler feet until the cabinet is level and the load is removed from all casters. Verify that the casters spin freely.
4. Tighten the locknuts on all four leveler feet.

Figure 3-5 Leveler Foot Adjustment



CXO-3829A

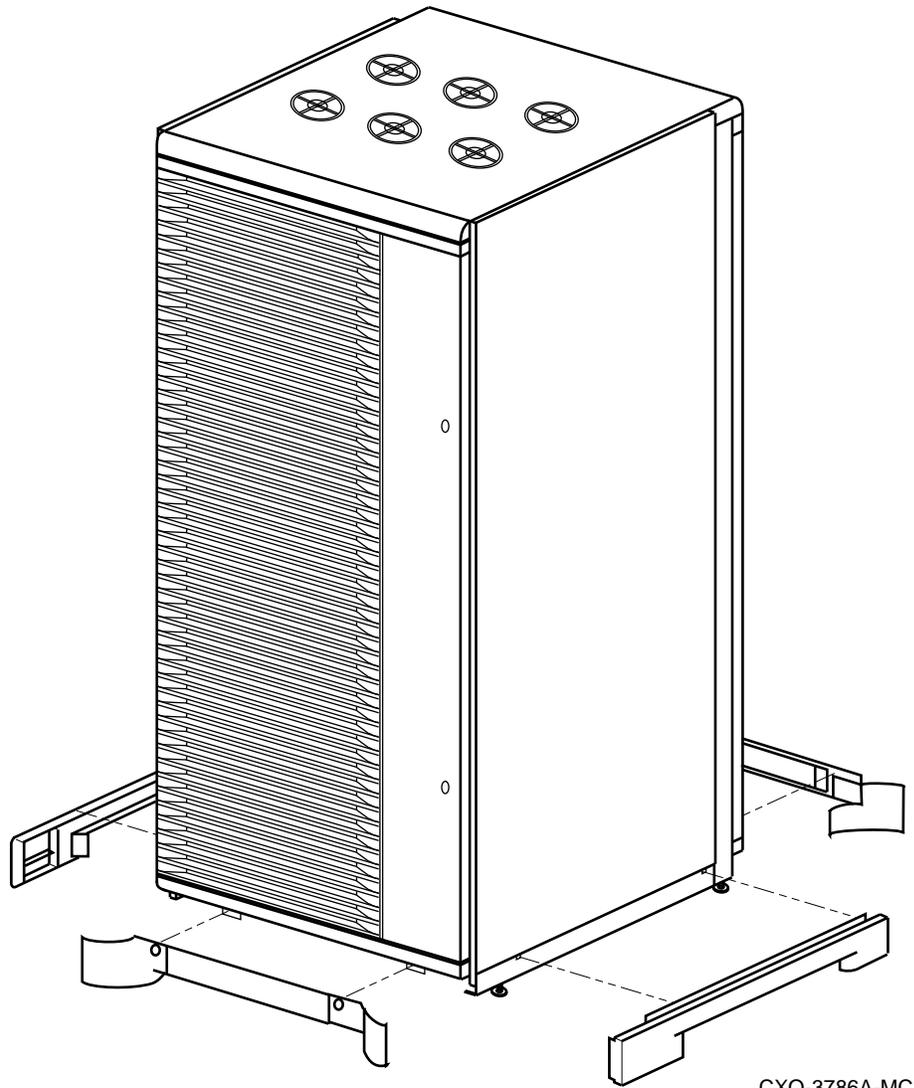
3.5 Installing the Skirt Kit

The skirt kit is packaged separately inside the corrugated carton with the cabinet. Installation of the skirt kit is optional. If you are installing the skirt kit, proceed as follows:

1. Unpack the skirt kit carton, using Figure 3-6 for reference, identify the right, left, front, and rear skirts.
2. Using Figure 3-6 for reference, position the skirts next to the cabinet.
3. The skirt fasteners are quarter-turn fasteners. Position each skirt against the cabinet such that the fasteners mate with the receptacles on the cabinet's base.
4. Using a screwdriver, push each fastener into its mating receptacle on the cabinet base, and turn it one quarter-turn clockwise.

(Once the skirts are installed, there is a small amount of play that allows them to be adjusted slightly up or down for proper alignment.)

Figure 3–6 Cabinet Skirt Installation



CXO-3786A-MC

3.6 Inspecting the Cabinet

Inspect the cabinet installation as follows:

1. Make sure that all four leveler feet are lowered to support the full weight of the cabinet and that the cabinet is level.
2. The remaining steps require access to the cabinet interior. To access the cabinet interior perform the following:
 - a. Release the door locks mounted on the smooth vertical panel of each door by turning the locks counterclockwise with a 5/32-inch hex wrench.
 - b. Swing the doors open.
 - c. Move the cabinet away from adjacent cabinets.

- d. Remove the side panels.

Note

There are three side panel hanger clips on each side of the cabinet. A matching set of hanger clips are attached to each side panel.

- e. Loosen the top cover by pushing up on its front and rear edges until it snaps free of its fasteners.

WARNING

The top cover is heavy and awkward to lift. Removing it requires two people. Failure to use sufficient personnel can result in personnel injury or equipment damage.

- f. Using two people, lift the top cover from the cabinet and set it aside.
 - g. Remove the bolts attaching the side panels to the top side rails of the cabinet.
 - h. Grasp a panel along its front and rear edges and lift up until the hanger clips disengage. Lift the panel away from the cabinet.
 - i. Repeat the previous step to remove the other panel.
3. Make sure that the circuit breaker on the CDUs is in the ○ (OFF) position.
 4. Make sure that all hardware within the cabinet is fastened securely, and that there are no loose pieces present in the cabinet interior.
 5. Check each of the cabinet fans to make sure that the blades turn freely and are not obstructed.
 6. Make sure that there are no obstructions to the airflow from the shelf blowers.
 7. Check the identification label on the rear of the cabinet to verify that the cabinet is configured to accept the power available at the site.
 8. Make sure that all ac power cords connected from the shelves and cabinet fans to the CDUs are firmly seated in their connectors at both ends.
 9. Make sure that all SBBs are seated firmly in their shelves.
 10. Make sure that any necessary external interface cables are installed and firmly seated in their connectors.
 11. Replace the side panels and top cover by reversing the procedures described in step 2.

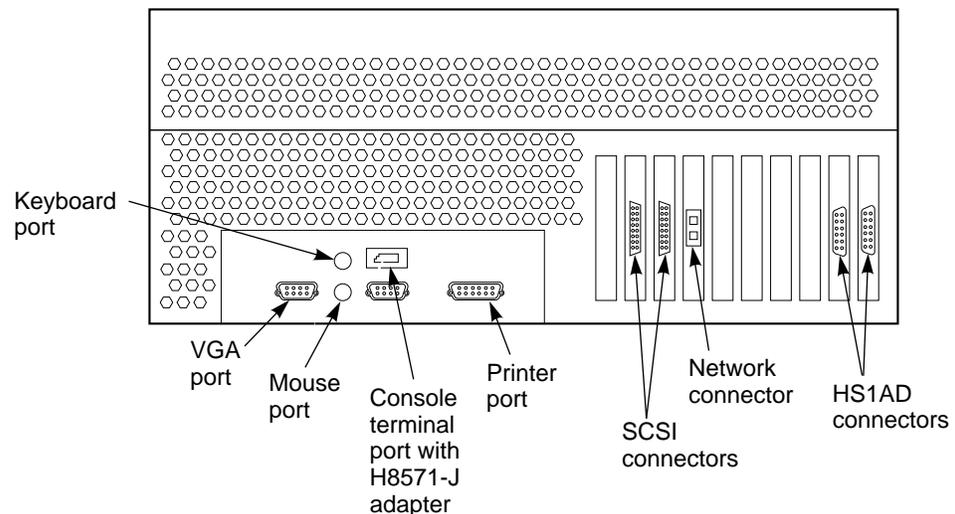
3.7 Connecting a Terminal to the StorageWorks Network Storage Array

Communicating with the StorageWorks Network Storage Array requires that you connect a terminal to it. Any terminal that supports ANSI control sequences can be used, including graphics displays that provide emulation of an ANSI compatible video terminal.

Use the following procedure to connect most EIA compatible terminals:

1. Make sure the power switch on the back of the terminal is OFF (0).
2. Connect one end of the terminal cable to the back of the terminal.
3. Connect the other end of the terminal cable to the EIA terminal port on the rear of the StorageWorks Network Storage Array cabinet using an H8571-J adapter as shown in Figure 3-7.

Figure 3-7 Connecting to the Terminal Port of the StorageWorks Network Storage Array



CXO-4774B-MC

4. Turn the terminal power switch to the ON position.
5. Set the terminal at 9600 baud, with 8 data bits, 1 stop bit, and no parity. Refer to your terminal documentation for terminal setup instructions.

3.8 Applying Power to the Cabinet

Once the cabinet has been inspected and the terminal has been connected to the StorageWorks Network Storage Array, apply power as follows:

1. Plug the primary power cable from CDU A into the appropriate site power receptacle.
2. Plug the auxiliary power cable from CDU B into the appropriate site power receptacle.
3. Switch the circuit breaker on each CDU to the | (ON) position.
4. Verify that all cabinet fans and shelf blowers are operating and that both status indicators on each shelf power supply SBB are illuminated. Refer to the *StorageWorks Solutions Shelf and SBB User's Guide* for further information on shelf status indicators.
5. Apply DC power to each server processor. The DC power switch is located on the OCP of the server processor. Figure 1-5 shows the location of the OCP.
6. Observe that the server processor power-on self-tests were successful by examining the terminal and looking for the self-test information as shown in Example 3-1.

Example 3-1 Server Processor Self Test Display

```
*** keyboard not plugged in...
ff.fe.fd.fc.fb.fa.f9.f8.f7.f6.f5.
ef.df.ee.f4.ed.ec.eb.....ea.e9.e8.e7.e6.e5.e4.e3.e2.e1.e0.
V1.1-1, built on Nov  4 1994 at 16:44:07
>>>
```

Customizing Network Array Operating System Parameters

This chapter describes the tasks that you perform to customize operating system parameters in the StorageWorks Network Storage Array.

4.1 Software Configuration Overview

Each StorageWorks Network Storage Array ships from the factory with the OpenVMS Alpha operating system preloaded on the shadowed system disk. You select network array processor operating system parameters the **first time** you boot the network array.

Because each server processor is a VMScluster node, it is important that you perform the software configuration in the recommended sequence. Table 4–1 describes the configuration tasks in this and following chapters. As you complete the setup procedures, THIS node is defined as the server processor in the network array that has the terminal connected to it and is running this process; OTHER node is the server processor without the terminal connected to it. N

Table 4–1 Software Configuration Tasks

Boot	THIS node	OTHER node	Refer to
First	Customize the parameters by running the Software Customization Procedure.	Powered off	Section 4.3, Executing the Software Customization Procedure.
Second	Configure the StorageWorks storage devices according to your needs.	Powered off	Chapter 5, Configuring the StorageWorks Network Storage Array Storage Devices.
Third	Connect the network array to the network interconnect.	Powered off	Section 8.1, Connecting to the Network Interconnect.
	Configure the network array's embedded cache.	Booted after the third boot of THIS node	Chapter 6, Managing the Network Storage Array's Read Cache.
	Configure the Network File System.		Chapter 7, Network File System.
	Ensure VMScluster system compatibility and network array readiness.		Chapter 8, Integrating the Network Array into the Network System.

Note

If you have an SWXNA-Ex/Jx (single server processor) model, complete the steps under THIS node and ignore the steps under OTHER node in Table 4-1.

Caution

Complete the first, second, and third boots on one server processor only. The setup information will be stored on the shadowed system disk for both server processors. Attempting to execute the Software Configuration Procedure on both server processor nodes simultaneously will result in corrupt data on the system disk.

4.2 VMSccluster System Considerations

Note

Section 4.2 only pertains to StorageWorks Network Storage Arrays which are to be used in a VMSccluster environment. Refer to Section 4.3 if your network array is to be installed into an environment other than a VMSccluster.

Introducing the StorageWorks Network Storage Array to a VMSccluster system requires you to consider some broad, clusterwide issues. Among them are the VMSccluster quorum scheme, enabling you to maintain an available VMSccluster system, and Volume Shadowing interoperability. This section addresses each of these topics with some guidelines for you to consider.

4.2.1 Considering the VMSccluster Quorum Scheme

The addition of a StorageWorks Network Storage Array to an existing VMSccluster may impact existing VMSccluster voting and quorum schemes. Two sample configurations may provide insight into possible VMSccluster voting and quorum schemes:

Example 1:

The network array is being added to an existing cluster with an existing and working quorum scheme. In this case, the network array is providing device serving only and does not contain a VMSccluster system disk for other cluster nodes. In this case, the network array should have ZERO votes and should contain no quorum disk (which may already exist in another part of the cluster). The expected votes on the network array should be set to the current value in the cluster.

Example 2:

The network array provides the VMScluster system disk for other VMScluster satellites (nodes without votes). There are no other voting nodes in the cluster. This case would occur in a new cluster being formed with a network array as the core. In this case, the nodes in the network array must provide all the votes for the cluster. For the two-node network arrays, a quorum disk on the network array is recommended so that the cluster can continue with only one server node operational. In the voting scheme, each of the two server nodes would have one vote; the quorum disk would have a single vote; and the expected votes for the cluster would be three.

Many other cluster configurations are possible, and appropriate quorum and voting schemes should be designed by knowledgeable system and network managers, perhaps with assistance from consulting services. Extensive information regarding quorum and voting issues may be found in the *VMScluster Systems for OpenVMS* manual and the *Guidelines for VMScluster Configurations* manual.

4.2.2 Volume Shadowing Interoperability

Each server processor in your network array is a node in your VMScluster running the OpenVMS Alpha Version 6.2 operating system. Introduction of this operating system version into your cluster has clusterwide implications that should be considered by knowledgeable system managers. OpenVMS Alpha Version 6.2 includes Phase II Volume Shadowing interoperability changes from previous ECO kits. If you are using Phase II Volume Shadowing, then you must install the Volume Shadowing interoperability patches on any cluster member running an operating system version earlier than OpenVMS Version 6.2 to ensure correct shadowing operation throughout your cluster.

More information regarding these patches is available in the Problem Description section of the release notes supplied with the patch kits. Please read these notes regarding the cluster impact of patch installation. For your convenience, the network array Version 2.0 CD-ROM contains the release notes and shadowing remedial kits current at the time the CD-ROM was built. The release notes and kits are located in the [SHADOW_KITS] directory.

Digital recommends that you contact the Digital Customer Support Center to obtain the latest Volume Shadowing ECOs for all operating system versions in your cluster. TIMA remedial kits containing these changes are available through Digital Customer Services worldwide.

4.3 Executing the Software Customization Procedure

Note

The Software Customization Procedure should be performed by an experienced system administrator who is knowledgeable about the installation site. Network knowledge also is required to complete the procedure.

The Software Customization Procedure (SCP) is self-explanatory and provides immediate help for determining responses to questions asked during the procedure. To get help, type a question mark (?) as your response to a procedure question and press the Return key.

The SCP requires specific information in order to complete successfully. Before starting, make sure that you have the information described in Table 4–2. (Table 4–2 can be used as a worksheet. Use the underlined space in column 1 to record your response to each of the questions in advance.) Defaults are provided for all parameters in the SCP except passwords.

Once completed, the SCP will not run again. However, if you abort the procedure before it completes, it will appear the next time you boot the system.

After completing the Software Customization Procedure, manage the StorageWorks Network Storage Array using normal OpenVMS Alpha system management procedures. For information on managing the network array, consult appropriate documentation on the OpenVMS Alpha operating system.

Table 4–2 Required Information for Customizing the OpenVMS Alpha Operating System

Your choice	Explanation
Current time and date	
_____	Current date
	Format: dd-mmm-yyyy
_____	Current time
	Format: hh:mm:ss
System Password	
_____	System Password
DECnet Node	
_____	DECnet node names of the StorageWorks Network Storage Array's server processors
_____	DECnet node addresses of the StorageWorks Network Storage Array's server processors

(continued on next page)

Table 4–2 (Cont.) Required Information for Customizing the OpenVMS Alpha Operating System

Your choice	Explanation
Cluster	
_____	Cluster Group Number
	Valid Ranges 1 - 4095, 61440 - 65535
_____	Cluster Password
Allocation Class Parameter	
_____	Value of the disk ALLOCLASS parameter.
	Valid Range 1 - 255
_____	Value of the tape ALLOCLASS parameter.
	Valid Range 1 - 255
Quorum Disk, if used	
_____	Name of the quorum disk
Note	
This disk cannot be a shadow set or a member of a shadow set.	

(continued on next page)

Table 4-2 (Cont.) Required Information for Customizing the OpenVMS Alpha Operating System

Your choice	Explanation
Vote Configuration	
_____	The number of votes each server processor will contribute toward quorum.
	Valid Range 0 - 127
_____	The expected number of votes in the cluster.
	Valid Range 1 - 127
Shadow Unit Number	
_____	The shadow unit number for the network array system disk.
	Valid Range 0 - 9999

To begin the Software Customization Procedure, type B at the >>> system prompt on THIS node, the one running the process. Do not run the procedure on the OTHER node, simultaneously. Examples 4-1 and 4-2 show sample Software Customization Procedure log files and provide sample answers to the questions in the Software Customization Procedure. All user input is shown in bold text except for passwords, which are not echoed, and carriage returns.

Example 4-1 Software Customization Procedure Log File—Part 1

```

OpenVMS (TM) Alpha Operating System, Version V6.2-1H3

%SHADOW-I-VOLPROC, DSA100: shadow master has changed. Dump file WILL be
written if system crashes. Volume Processing in progress.

Copyright (c) 1996 Digital Equipment Corporation. All rights reserved.

+-----+
|           Welcome to the
|           StorageWorks (TM) Network Storage Array
|           Software Customization Procedure
+-----+

This Network Storage Array has two server processors.

```

(continued on next page)

Example 4-1 (Cont.) Software Customization Procedure Log File—Part 1

W A R N I N G

The Software Customization Procedure (SCP) must be invoked from only one of the server processors. Executing the SCP simultaneously on both server processors will result in data corruption on the network array system disk.

If the other server processor is currently executing or has executed the SCP, you must power off this server processor now. This server processor can then be powered on and booted after the other server processor has completed its third boot.

Press Return to continue:

Please wait while the Software Customization Procedure initializes the system. This will take less than one minute.

```
+-----+
|                                     |
|               Software Customization Procedure               |
|                   Introduction                               |
|                                     |
+-----+
```

The Software Customization Procedure will ask you to enter the information required for the factory installed software to be custom configured for your site.

To answer a question, type the requested information and press Return. To get help at any time, type a question mark (?) and press Return.

Press Return to continue:

No OpenVMS licenses have been registered on this system.

To register your license PAKs, type the following command after the system restarts and you have logged in to the system manager's account:

```
$ @SYS$UPDATE:VMSLICENSE
```

Press Return to continue:

The following information is essential to customize the network array.

You cannot complete the Software Customization Procedure if you do not have all the essential information. If you do not have this information, you can obtain it from your system manager or network administrator.

Default values for the cluster parameters are provided. These values can be taken if this network array is installed on a network which does not have existing DECnet nodes or VMScluster systems.

- o The current date and time
- o The password for the system manager's account on the network array
- o Cluster Information
 - The DECnet node names and addresses
 - The cluster group number and password
 - The value of the allocation class parameter for disk and tape
 - The name of a quorum disk if a quorum disk will be installed on the network array (this disk cannot be a shadow set or a member of a shadow set)
 - The number of votes each server node will contribute toward quorum
 - The number of expected votes in the cluster
 - The shadow unit number for the network array system disk

Are you ready to proceed with the Software Customization? (Y/N) [YES]

At the end of the Software Customization Procedure, the server processor will restart itself. Then it will be ready for use.

(continued on next page)

Example 4-2 Software Customization Procedure Log File—Part 2

Waiting for the cluster configuration to complete...

\ 21-AUG-1996 11:13:41 /

Cluster configuration complete.

```
+-----+
|                               |
|                               |
+-----+
```

The installation of the Network File Serving software will take about 2 minutes. Please wait...

The Network File Serving software has been successfully installed.

After the network array has formed a cluster (or joined an existing cluster) and the SYSUAF, RIGHTS LIST, and license database are defined, please execute the following command to complete NFS installation tasks:

```
$ @SYSS$MANAGER:NSA$SCP_POST_CONFIG.COM
```

Press Return to continue:

```
+-----+
|                               |
|                               |
+-----+
```

A customized version of PerfectCache (TM) by Raxco Software, Inc. is being installed on the network array.

This installation will take less than 1 minute. Please wait...

PerfectCache has been successfully installed. Your PerfectCache license has been activated.

Included in the software already installed on this StorageWorks Network Storage Array is a derivative of PerfectCache, which is a copyrighted product and which contains proprietary and confidential information. Copyright Raxco Software, Inc. 1996. That product may be used only on this Network Storage Array, and may not be reproduced, transferred to any other processing unit, reverse engineered, decompiled or provided to any other person without the express prior consent of Raxco Software, Inc.

Press Return to continue:

```
+-----+
|                               |
|                               |
+-----+
```

This is the end of the data entry part of the Software Customization Procedure. You do not need to provide any more information.

The server processor will now shut down and restart automatically as a standalone node; it will not join or create a cluster. After this reboot, you can login to the system manager's account using the user name SYSTEM and the password that you specified earlier. You may then install any license PAK's, configure the storage attached to the network array, add a command to the startup procedure to mount the quorum disk (if any), perform additional system customization, etc.

(continued on next page)

Example 4–2 (Cont.) Software Customization Procedure Log File—Part 2

The third time the server processor is started, it will attempt to form or join a VMScluster using the parameters you have entered above and will be accessible to network clients. Do not boot the other server processor until this third boot is completed. After both server processors have joined the VMScluster, you can execute the post configuration procedure: \$ @SYS\$MANAGER:NSA\$SCP_POST_CONFIG.COM.

Press Return to complete the Software Customization Procedure:

The server processor will now restart automatically. Please wait...

Pressing the Return key at this point completes the Software Customization Procedure and results in the system prompting you for a user name and password. Continue on to Chapter 5 to configure the storage devices.

Note

If you specified a quorum disk as part of the StorageWorks Network Storage Array software installation, the quorum disk must be mounted after the OpenVMS Alpha operating system is running. As soon as possible, add a command to mount the quorum disk on the network array to the startup procedure for the server's processor. This allows you to maintain quorum in the event one of the server processors fails (see Section 9.1.1).

Configuring the StorageWorks Network Storage Array Storage Devices

The StorageWorks Network Storage Array requires you to configure the StorageWorks storage system to meet your particular storage needs. The *StorageWorks Array Controllers HS Family of Array Controllers User's Guide* provides a detailed description of stripesets and RAIDsets that you can configure for the StorageWorks Network Storage Array.

This chapter begins by providing a few general guidelines to help you decide how you may want to configure the StorageWorks storage system, continues with the steps required to use the CFMENU utility from a terminal connected to the server processor, and concludes by providing a sample configuration.

5.1 General Considerations for Configuring a StorageWorks Storage System

Configuring the storage system is the process of defining, through parameters, the logical organization of the storage devices. The storage devices may be configured as one or more of the following:

- **Mirrorset**—Two member disks containing identical images of user data.
- **Stripeset**—A virtual disk drive with its physical data spread across multiple physical disks.
- **RAIDset**—Three or more physical disks that are connected to present an array of disks as a single virtual unit to the host.
- **Spareset**—A pool of disk drives used by the device channel processor to replace failing members of a RAIDset.
- **Failedset**—A group of disk drives that have been removed from RAIDsets due to a failure or a manual removal.
- **Container**—An entity that is capable of storing data, whether it is one physical device or a group of physical devices. A disk, a stripeset, and a RAIDset are examples of a container.
- **Passthrough Container**—A virtual device used to pass SCSI commands that perform operations other than simple device read or write operations. A passthrough container typically is used with tape libraries that contain standard tape devices plus a robotic mechanism for selecting and loading cartridges from a pool of tapes. A passthrough container must exist to pass SCSI commands to the robotic mechanism.
- **Unit**—A logical entity composed of one or more devices and treated as a single addressable storage structure by the cluster members. Units can be single devices, stripesets, or RAIDsets.

5.2 Running CFMENU from the StorageWorks Network Storage Array Terminal

Before CFMENU can be invoked from a terminal connected to the StorageWorks Network Storage Array, you must establish a logical connection to the device channel processor. To create a logical connection to the device channel processor, perform the following steps:

1. Log into the system manager's or other appropriate account on the StorageWorks Network Storage Array (the account must have DIAGNOSE privilege).
2. At the DCL prompt, enter the SHOW CLUSTER command. The system will respond with default cluster information similar to that shown in Figure 5-1 and return control to the DCL command level.

Figure 5-1 Example Cluster Information

View of Cluster from Node: NSTAR1

20-OCT-1996 10:51:09

SYSTEMS		MEMBERS
NODE	SOFTWARE	STATUS
NSTAR1	VMS V6.2	MEMBER
NSTAR2	VMS V6.2	MEMBER
NODE01	VMS V6.2	MEMBER
NODE02	VMS V6.2	MEMBER
NODE03	VMS V6.2	MEMBER
NODE04	VMS V6.2	MEMBER
HS1CP1	HSD Vxx	
HS1CP2	HSD Vxx	
HS1CP3	HSD Vxx	
HS1CP4	HSD Vxx	
HS1CP5	HSD Vxx	
HS1CP6	HSD Vxx	
HS1CP7	HSD Vxx	
HS1CP8	HSD Vxx	

CXO-5430A-MC

3. Identify the names of the device channel processors (note, there may be more than one). The device channel processor can be identified by searching for the nodes that are named HS1CP1, HS1CP2, and so on.
4. Change the terminal format to accommodate 132 characters with this command:

```
$ SET TERMINAL/WIDTH=132
```

5. Logically connect the StorageWorks Network Storage Array terminal to the device channel processor by entering the following command at the DCL prompt:

```
$ SET HOST/DUP/SERVER=MSCP$DUP/TASK=CLI node-name
```

Where:

node-name is the name assigned to the device channel processor (for example, HS1CP1).

The device channel processor responds with its prompt (for example, HS1CP>).

- Enter the following command to start the CFMENU utility:

```
HS1CP> RUN CFMENU
```

The CFMENU utility begins executing. After CFMENU's initial message, press the Return key to get to CFMENU's main menu as shown in Figure 5–2. Any recognized devices that have not been added to your configuration will appear, with their Port Target LUN (PTL), in the column to the right of the menu options. All the devices shown in Figure 5–2 are attached but not yet configured.

Figure 5–2 CFMENU Main Menu

MAIN MENU:	Unconfig'd Dev.PTLs	CFMENU Config'd PTLs	Configuration Device Name	Menu Product ID	Utility	Stor.set Name	Stor.set Typ/Sz	Chnk Size	Trn sp.	In- it'd	Re- duc	W Unit	W P	B
1. Add/delete devices														
2. Add/delete mirrorsets														
3. Add/delete stripesets		disks: 100	DISK100	RZ35	(C) DEC					N	Y	D100	N	Y
4. Add/delete raidsets/ sparesets/failedsets		110	DISK110	RZ28	(C) DEC					N	Y	D110	N	Y
5. Add/delete passthrough		120	DISK120	RZ26L	(C) DEC					N	Y	D120	N	N
6. Initialize devices and/or storagesets		150	DISK150	RZ73	(C) DEC					N	Y	D150	N	N
7. Add/delete units		210	DISK210	RZ35	(C) DEC					N	Y	D210	N	N
8. Setup terminal		220	DISK220	RZ26	(C) DEC					N	Y	D220	N	N
9. Exit CFMENU		240	DISK240	RZ74	(C) DEC					N	Y	D240	N	N
		300	DISK300	RZ28	(C) DEC					N	Y	D300	N	Y
		320	DISK320	RZ26L	(C) DEC					N	Y	D320	N	N
		340	DISK340	RZ29B	(C) DEC					N	Y	D340	N	N
		420	DISK420	RZ74	(C) DEC					N	Y	D420	N	Y
		440	DISK440	RZ74	(C) DEC					N	Y	D440	N	N

CXO-5279A-PH.EPS

Note

If your terminal is capable of displaying more than 24 rows, you may want to enter option 8 from the main menu to set the number of rows CFMENU will display.

Table 5–1 describes the information headings of the main menu. To avoid confusion, the information headings are presented exactly as they appear on the main menu.

Table 5–1 CFMENU Information Headings

Information Heading	Description
Main Menu	Lists the major operations that are available.
Unconfig'd Dev.PTLs	Unconfigured PTLs—The device PTLs that are recognized by the HS1CP device channel processor but that have not yet been added to the configuration.
Config'd PTLs	Configured PTLs—The device PTLs that have been configured by the HS1CP device channel processor.
Device Name	Device name—The name automatically assigned to the device when it is configured by the HS1CP device channel processor. The name can be changed through use of the CLI commands.
Product ID	Product ID—Information that identifies the device.
Stor.set Name	Storage set name—The name assigned by CFMENU to the storage set. The name is assigned using one of the following conventions: <ul style="list-style-type: none">• Sx for stripesets• Rx for RAIDsets• Px for passthrough containers The number denoted by <i>x</i> is assigned sequentially beginning with 1 for each type of storage set. The name can be changed through use of the CLI commands. (See the <i>StorageWorks Array Controllers HS Family of Array Controllers User's Guide</i> for complete descriptions of the CLI commands.)
Stor.set Type	Storage Set—The type of storage set: STRP for stripesets, RAID for RAIDsets, PASS for passthrough containers.
Chnk Size	Chunksize—The size, in blocks, of data transfers with the device.
Trn sp.	Transportable—The letter, Y or N, to indicate if the device is transportable or not. A transportable device is one that does not have metadata written on it. A nontransportable device has a small amount of metadata written on it.
In-it'd	Initialized—The letter, Y or N, to indicate if the device has been initialized.
Re-duc	Reduced—Valid for RAIDsets only, indicates the RAIDset is missing one member.
Unit	Unit—The logical number of the unit as assigned by you and preceded by the letter D for disk or T for tape.
W P	Write Protected—The letter, Y or N, to indicate whether the device is write protected.
W B	Write-back—The letter, Y or N, to indicate if the device has write-back caching enabled.

5.3 Configuring Storage Devices Using CFMENU

5.3.1 Using CFMENU

Configuring a storageset with CFMENU is as easy as choosing menu options and responding to prompts. Just remember to move through the Main Menu from top to bottom for each storageset you want to configure.

Note

Press "D" to scroll down CFMENU's columns; press "U" to scroll up.

Complete the following steps to configure a storageset.

1. Add the disk drives that you want to combine into a storageset.
2. Add the storageset to the controller's list of available storagesets. CFMENU prompts you for a storageset number and the switches you want to apply to the storageset.
3. Initialize the storageset. CFMENU prompts you for the Initialize switches that you want to apply to it.
4. Assign a unit-number to the storageset. CFMENU prompts you the the Unit switches that you want to apply to it.
5. Exit CFMENU or repeat these steps to configure another storageset.

5.3.2 Considerations for Using CFMENU

- CFMENU automatically names each storageset according to a simple naming convention: *Mn* for mirrorsets, *Sn* for stripesets and *Rn* for RAIDsets, where *n* is a sequentially indexed number. You will have to configure your storagesets manually if you want to use your own naming scheme.
- You only need to run CFMENU on one device channel processor. The OTHER device channel processor automatically inherits the configuration you create with CFMENU as long as the device channel processors are powered up and in dual redundant configuration.

5.4 Modifying Allocation Classes

Note

If you selected the default parameters during the Software Customization Procedure, completing Section 5.4, Modifying Allocation Classes is not necessary. If you supplied your own values for Disk and Tape Allocation, complete the following procedure.

During the Software Customization Procedure, you were prompted to type in the MSCP and TMSCP allocation classes of the network array. The HS1CP device channel processors must have the same MSCP and TMSCP allocation class as the server processors. Set the allocation classes with this procedure:

1. From the currently active menu, type in the number to select the option:

Return to main menu

2. Exit from the CFMENU utility to return the the HS1CP> prompt. From the Main Menu type in 9 at the Enter Menu choice:

```
(1,9) [9]?
```

3. Enter the following command to determine the MSCP allocation class and the TMSCP allocation class:

```
HS1CP> SHOW THIS
```

4. If the display shows the correct allocation class information, then proceed with step 7. If the display shows the incorrect or incomplete allocation class information, then continue.

5. Set the MSCP allocation class by entering this command:

```
HS1CP> SET THIS MSCP_ALLOCATION_CLASS = n
```

Where:

n is the number you entered for the network array.

The following will be displayed on the terminal screen:

```
Warning 4020: A restart of both this and the other controller is
              required before all the parameters modified will
              take effect.
```

```
Restart of this controller required
```

```
Restart of the other controller required
```

6. Set the TMSCP allocation class by entering this command:

```
HS1CP> SET THIS TMSCP_ALLOCATION_CLASS = n
```

Where:

n is the number you entered for the network array.

The following will be displayed on the terminal screen:

```
Warning 4020: A restart of both this and the other controller is
              required before all the parameters modified will
              take effect.
```

```
Restart of this controller required
```

```
Restart of the other controller required
```

7. Restart the other HS1CP device channel processor with this command:

```
HS1CP> RESTART OTHER
```

These messages will appear on the screen:

```
Received LAST GASP message from other controller, Fail Code: 08080000
```

```
HS1CP>
```

```
%PAx0, Software is Closing Virtual Circuit -- REMOTE NODE HS1CPx
```

8. Restart the HS1CP device channel processor to which you are currently connected with this command:

```
HS1CP> RESTART THIS
```

These messages will appear on the screen:

```
%PAx0, Software is Closing Virtual Circuit -- REMOTE NODE HS1CPx
```

```
%HSCPAD-F-NOLOCEXE, Local program not executing
```

```
-SYSTEM-F-VCBROKEN, virtual circuit broken
```

```
%HSCPAD-S-END, Control returned to node XXXXXX
```

At this time you will return to the DCL prompt from which you connected to the HS1CP, and your configuration procedure is complete.

5.4.1 Configuring Storage for Additional Device Channel Processor Pairs

The SWXNA-Bx/Gx StorageWorks Network Storage Array contains two pairs of dual-redundant device channel processors, and four pairs when coupled to an SWXNA-Cx expansion cabinet. Configure the storage to be attached to each dual-redundant pair as described in Section 5.2 through Section 5.4.

5.5 Saving the Configuration Setup

After defining each Device Channel Processor or dual-redundant pair configuration setup, make sure it is printed and kept available to assist in servicing the subsystem in the future. Make a new printout each time you change your configuration.

To capture your device setup, you should perform the following steps:

1. Enter the following command at the DCL prompt to logically connect the StorageWorks Network Storage Array terminal to the device channel processor and record the results of the session in a file called CONFIG.LOG in the SYSSMANAGER directory on the network array's system disk.

```
$ SET HOST/DUP/SERVER=MSCP$DUP/TASK=CLI/LOG=SYSSMANAGER:CONFIG.LOG node-name
```

Where:

node-name is the name assigned to the device channel processor (for example, HS1CP1).

The device channel processor responds with a brief display and its prompt.

```
Copyright (C) Digital Equipment Corporation 1994
HS1CP Firmware version E35D-0, Hardware version AX01
Last fail code: 018000A0
Press " ?" at any time for help.
HS1CP>
```

2. At the HS1CP> prompt enter the following command:

```
HS1CP> SHO DEVICE FULL
```

The following information will be captured in the CONFIG.LOG file in the SYSSMANAGER directory on the network arrays's system disk as it appears on the screen:

Name	Type	Port	Targ	Lun	Used by
DISK100	disk	1	0	0	D1100
	DEC RZ28	(C)	DEC	D41C	
	Switches: NOTTRANSPORTABLE				
DISK110	disk	1	1	0	D1110
	DEC RZ28B	(C)	DEC	0003	
	Switches: NOTTRANSPORTABLE				
DISK120	disk	1	2	0	D1120

```
DEC      RZ28      (C) DEC D41C
Switches:
NOTTRANSPORTABLE
DISK130   disk           1   3   0           D1130
DEC      RZ28      (C) DEC D41C
Switches:
NOTTRANSPORTABLE
```

3. Exit the display by entering this command at the prompt:

```
HS1CP> EXIT
```

At this point, the HS1CPs and their storage devices have been configured. You should leave the terminal attached to the server processor while you perform other tasks in the installation procedure.

Managing the Network Storage Array's Read Cache

Caching is the use of memory to hold frequently accessed disk data so that physical disk I/Os can be avoided. If caching software is running, your system can complete user disk I/O read requests immediately rather than waiting for the disk to transfer data. Caching significantly increases the I/O access speed for read requests.

PerfectCache® by Raxco Software, Inc. is embedded in the StorageWorks Network Storage Array's operating software and is automatically loaded, configured, and started at the time the operating system is loaded. By default, all available disks are cached.

Note

Included in the software already installed on this StorageWorks Network Storage Array is a derivative of PerfectCache, which is a copyrighted product and contains proprietary and confidential information. Copyright Raxco Software, Inc. 1996. That product may be used only on this network storage array, and may not be reproduced, transferred to any other processing unit, reverse engineered, decompiled or provided to any other person without the express prior written consent of Raxco Software, Inc.

PerfectCache will cache drives only that are local to the network array. The disks must be mounted on the network array, Files-11, and shared (/system or /cluster).

PerfectCache will not cache host-based RAIDsets, stripesets, or shadowsets. Controller-based RAIDsets, stripesets, and shadowsets are not affected by this restriction.

You can safely add additional StorageWorks Storage Network Arrays running their embedded PerfectCache into the cluster. However, no node in the VMScLuster can run any host-based, block-served caching software other than the PerfectCache that is shipped with network arrays. OpenVMS Virtual I/O Cache (VIOC) can be run because it is a file cache, not a block cache. HS1CP cache can be run, because it is not host-based.

The standard version of PerfectCache can be installed provided the embedded version is stopped on the network array as described in Disabling PerfectCache Startup and the network array is rebooted.

Caution

Data corruption could result if another host-based cache system is installed on the network array or on other clients in the network.

6.1 Using the PerfectCache Menu

The PerfectCache Menu allows you to access various features and tools of PerfectCache. Enter the following command to access the PerfectCache Menu:

```
$ @PERFECT_CACHE_MENU
```

After displaying the copyright screen, the following menu is displayed:

PerfectCache/Alpha

- 1) Manage Cacher
- 2) Display Caching Summary
- 3) Display Caching Summary Continuously
- 4) Reset Caching Statistics
- 0) Exit

Note: The cacher is currently running.

Enter option number [0]:

6.1.1 Manage Cacher

Entering option number 1 on the PerfectCache Menu displays the Manage Cacher screen as follows:

```
Cache Manager on node NSTAR1
Screen 1 of 1
NSTAR1: Cacher Running 1 NSTAR2:: Cacher Running 1
```

Device	size/ %free	Local Mount	Caching Act	Req	Device	size/ %free	Local Mount	Caching Act	Req
1) 2) \$9\$DKA0: 3)	2.1 4)	N 5)	N 6)	N 7)	10) \$9\$DUA150:	2.0/76	Y	Y	Y
2) \$9\$DKA100:	2.1	N	N	N	11) \$9\$DUA210:	0.9/98	Y	Y	Y
3) \$9\$DKA400:		N	N	Y	12) \$9\$DUA220:		N	N	Y
4) \$9\$DKA500:		N	N	Y	13) \$9\$DUA240:	3.6/83	Y	Y	Y
5) \$9\$DKB300:		N	N	Y	14) \$9\$DUA300:	2.1/55	Y	Y	Y
6) \$9\$DKB500:	1.6/69	Y	Y	Y	15) \$9\$DUA320:		N	N	Y
7) \$9\$DUA100:	0.9/48	Y	Y	Y	16) \$9\$DUA340:	4.3/26	Y	Y	Y
8) \$9\$DUA110:	2.1/10	Y	Y	Y	17) \$9\$DUA420:	3.6/38	Y	Y	Y
9) \$9\$DUA120:		N	N	Y	18) \$9\$DUA440:		N	N	Y

Option (Exit, Quit, Next, Previous, Disk, Cacher, Help)

Note

While the menu is displayed, you can access help by pressing the H or ? keys.

The following list describes the information included in the Cache Manager Screen.

- ❶ The Cacher is running on both server processors, NSTAR1 and NSTAR2.
- ❷ Each disk is assigned a reference number for making PerfectCache changes.
- ❸ Device—The name of the disk.
- ❹ size/%freeThe size of the disk in gigabytes and the percentage of that disk which is available. No data in this field indicates the disk is not mounted or not viewable by PerfectCache for some other reason.
- ❺ Local Mount—Y indicates the disk is mounted on the local server processor.
- ❻ Caching Act—Caching active. Y indicates this disk is being cached by PerfectCache.
- ❼ Caching Req—Caching requested. Y indicates PerfectCache has been requested to cache this disk.

Normally Caching Active and Caching Requested should be the same. The following are some of the reasons that a disk may be requested to cache, but is not:

- Disk is not mounted locally.
- Disk is mounted either 'foreign' or privately.
- Disk is a member of a shadow set.

6.1.1.1 Using the Cache Manager Screen

The Cache Manager screen allows you to start and stop the cacher on the Network Storage Array or to start and stop caching individual disks. The following list summarizes the available options for the Cache Manager screen:

- E - EXIT the screen and effect the selected changes.
- Q - QUIT and ignore any changes.
- N - display the NEXT screen.
- P - display the PREVIOUS screen.
- D - change the DISK caching request.
- C - turn the CACHER on and off.
- H - display HELP.

Start/Stop the Cacher for the Network Array

Pressing the C key displays a message asking you to confirm starting or stopping the cacher on the network array. The change takes effect on *both* server processors when exiting (pressing the E key) the screen and also enables or disables PerfectCache startup for the next boot. Quitting (pressing the Q key) the screen ignores all changes made during the session.

Note

Once stopped, caching will remain off until you restart it using the Cache Manager screen.

Start/Stop the Cacher for individual disks

Pressing the P (Previous) and N (Next) keys switches among the Cache Manager screens to list all disk drives available on the StorageWorks Network Storage Array. Pressing the D key queries you for the disk reference number. Entering the disk reference number toggles the Caching Requested column entry for the selected disk. The change takes effect on both server processors when exiting (pressing the E key) the screen. Quitting (pressing the Q key) the screen ignores all changes made during the session.

Note

Once caching is stopped on a disk, caching will remain off until you restart it using the Cache Manager screen.

6.1.2 Display Caching Summary

Entering option number 2 on the PerfectCache Menu displays statistics about the caching activity since the Cacher was started on a disk-by-disk and total basis. The Cacher must be running to use this option. The following is an example of a summary screen:

```
Copyright (C) 1996          PerfectCache/AXP Summary          6-MAY-1996 13:51:27
Raxco, Inc.                Node: NSTAR1                          Mode: Total

Device      Reads      Read      Hit      Blocks      Blocks      Cache
              Hits      Rate      Read      Hit          Pages
$9$DKB500:      0          0      0.00      0          0          0
$9$DUA100: 278604    269359  96.68    1135900    547625     0
$9$DUA110:      0          0      0.00      0          0          0
$9$DUA210:      0          0      0.00      0          0          0
$9$DUA240:      0          0      0.00      0          0          0
$9$DUA300:      0          0      0.00      0          0          0
$9$DUA340:      0          0      0.00      0          0          0
$9$DUA420:      0          0      0.00      0          0          0
$9$DUA420:      43         3      6.98     2062        30         0
$9$DUA150:      0          0      0.00      0          0          0
Total:        278755    269396  96.64    1141064    548323     2
```

Press Return to continue:

Statistics summarized since caching started or the statistics were reset are:

- Reads—Number of read requests
- Read Hits—Number of read requests satisfied the cache without actually performing a transfer from the disk
- Hit Rate—Percentage of read requests satisfied from cache
- Blocks Read—Total number of blocks read
- Blocks Hit—Number of blocks read from cache
- Cache Pages—Number of 8K byte pages of memory used to hold cache data for each disk

6.1.3 Display Caching Summary Continuously

Entering option number 3 on the PerfectCache Menu uses the same display as Section 6.1.2 to display current cache activity. The screen is updated every 5 seconds to show the statistics on a disk-by-disk and total basis. The Cacher must be running to use this option.

6.1.4 Reset Caching Statistics

Entering option number 4 on the PerfectCache Menu resets the information for the Caching Summary (Section 6.1.2) which displays caching statistics for all devices since the Cacher process was started.

The Cacher must be running to reset the statistics.

6.2 PerfectCache Commands

PerfectCache is automatically loaded and started when the network array is booted. As an alternative to using the PerfectCache menu, commands may be issued from a terminal to:

- Disable PerfectCache from starting on the next boot
- Enable PerfectCache to start on the next boot
- Show the current state of the PerfectCache startup

Refer to Section 3.7, Connecting a Terminal to the StorageWorks Network Storage Array for the procedure for connecting a terminal to the server processor in order to issue PerfectCache commands.

Disabling PerfectCache Startup

If you do not want PerfectCache to automatically start on your system, type the following command:

```
$ MCR SYSMAN STARTUP DISABLE FILE NSA$PCS_STARTUP.COM
```

This command will take effect the next time the network array is booted.

Enabling PerfectCache Startup

This command needs to be issued only if PerfectCache had been previously disabled and you now want it enabled. To enable PerfectCache, enter the following command:

```
$ MCR SYSMAN STARTUP ENABLE FILE NSA$PCS_STARTUP.COM
```

PerfectCache will now start after the next network array boot.

Displaying PerfectCache Current Startup State

To display the current startup state of PerfectCache, enter the following command:

```
$ MCR SYSMAN STARTUP SHOW FILE NSA$PCS_STARTUP.COM /FULL
```

If PerfectCache is *enabled*, you will see a message similar to the following:

```
%SYSMAN-I-COMFIL, contents of component database on node NSTAR1
Phase      Mode   File
-----
LPMAIN     DIRECT NSA$PCS_STARTUP.COM
           Enabled on All Nodes
           P1:<def> P2:<def> P3:<def> P4:<def> P5:<def> P6:<def> P7:<def> P8:<def>
```

If PerfectCache is *disabled*, you will see a message similar to the following:

```
%SYSMAN-I-COMFIL, contents of component database on node NSTAR1
Phase      Mode   File
-----
LPMAIN     DIRECT NSA$PCS_STARTUP.COM
           Disabled on All Nodes
           P1:<def> P2:<def> P3:<def> P4:<def> P5:<def> P6:<def> P7:<def> P8:<def>
```

Network File System

7.1 Network File System

The Network File System (NFS™) is a protocol developed by Sun Microsystems, Inc. that uses Internet Protocol (IP) to allow computers to access each other's files as if they were local. When applications request file operations, such as open, store, or retrieve a file, the operating system passes the request to either the local file system or the NFS Client. When it receives a request, the NFS client contacts the appropriate NFS Server on the remote machine which performs the requested operation.

NFS provides users with the following capabilities:

- "Transparent" user access to remote files.
- Inter-host access to remote files without moving them physically across the network.

7.1.1 NFS Server

The NFS Server enables remote users to access files that physically reside on the Network Storage Array running the TCP/IP Services for OpenVMS (UCX) software. Remote users can do this from any host that supports the NFS Client. UNIX supports only sequential byte-stream file formats. OpenVMS on the network array supports index, relative, and sequential files that have various record formats. The UCX implementation of NFS allows access to the following record formats:

- Fixed length
- Variable length
- Variable with fixed-length control (VFC)
- Stream (including STREAM_LF and STREAM_CR)
- Undefined

7.1.2 NFS Client

The NFS Client accepts file operation requests from the operating system and contacts the NFS Server on the network storage array, which performs the requested operation.

7.1.3 PC-NFS Client Daemon

NFS client software can also run on a personal computer (PC). A PC using NFS client can mount a remote UNIX®, OpenVMS, or mainframe disk volume and access files as though they reside locally on the PC.

Network security in the NFS mode is accomplished by identifying every user with a User Identification (UID) / Group Identification (GID) pair. Personal computers lacking this information, must request their UID / GID from a remote server. They obtain this information from the UCX Proxy Database. With this information, PC users can then mount remote NFS file systems.

7.2 Checklist for Configuring the NFS Server

Complete the following worksheet to prepare to configure NFS on the network array. You will be asked various questions during the configuration about THIS node and OTHER node. THIS node is defined as the server processor in the network array, which has the terminal connected to it and is running this process. OTHER node is the server processor without the terminal connected to it.

Note

For the NSA's NFS to work properly with Sun workstations, the following line must be added to the file `/etc/init.d/nfs.server` on the Sun workstation as the first line of the file:

```
/usr/sbin/ndd -set /dev/ip ip_path_mtu_discovery 0
```

Complete the following worksheet to prepare to configure the NFS Server on the network array.

Table 7-1 Checklist for Configuring the NFS Server

Your Choice	Explanation
-------------	-------------

General Questions

_____ For security purposes the Proxy Database is created with no world access. This denies nonprivileged users the use of commands which access the Proxy Database. Do you want to grant world read access to nonprivileged users? Yes or No

The default is no.

_____ Select your geographic area from the following chart.

```
Selection  Geographic Area
-----
u      USA
ea     Eastern Australia
ca     Central Australia
wa     Western Australia
ee     Eastern Europe
ce     Central Europe
we     Western Europe
g      Greenwich
n      none of the above
-----
```

_____ Does your area alternate between Daylight Savings and Standard time? Yes or No

_____ What is the time zone for your area?

Example: Greenwich or Eastern

Domain Configuration

_____ What is your Internet domain?

In fully qualified host name 'myclient.mytown.dec.com',
the Internet domain is mytown.dec.com

(continued on next page)

Table 7-1 (Cont.) Checklist for Configuring the NFS Server

Your Choice	Explanation
Interface Configuration	
<i>Information about this section can be found in the Networking and TCP/IP section of the DEC TCP/IP Services for OpenVMS, Concepts and Planning document.</i>	
_____	What is the Internet address for THIS node?
	Example: TOPCPU: 16.253.32.112
_____	What is the Internet address for OTHER node?
	Example: BOTCPU: 16.253.32.111
_____	What is the Internet network mask for THIS node and OTHER node?
	Example: [255.0.0.0]:
_____	What is the broadcast mask for THIS node and OTHER node?
	Example: [16.255.255.255]:

(continued on next page)

Table 7-1 (Cont.) Checklist for Configuring the NFS Server

Your Choice	Explanation
Cluster Host Configuration	
_____	<p>What is the unqualified Cluster Host name:</p> <p>Unqualified host name example: In fully qualified host name 'myclient.mytown.dec.com', the unqualified host name is myclient</p>
_____	<p>What is the Internet address for the Cluster Host?</p> <p>Example: 16.253.32.115</p>
_____	<p>What is the Internet network mask for the Cluster Host?</p> <p>Example: [255.0.0.0]</p>
_____	<p>What is the broadcast mask for the Cluster Host?</p> <p>Example: [16.255.255.255]</p>
Portmapper Server Configuration	
_____	<p>Approximately how many NFS Clients will be served?</p> <p style="text-align: right;">(continued on next page)</p>

Table 7-1 (Cont.) Checklist for Configuring the NFS Server

Your Choice	Explanation
Client Host and Proxy Configuration	
Information about this section can be found in the <i>DEC TCP/IP Services for OpenVMS, Installation and Configuration</i> and <i>DEC TCP/IP Services for OpenVMS, Management</i> documents.	
For each client host ...	
_____	What is the fully qualified name of a client host? Fully qualified client host name 'myclient.mytown.dec.com',
_____	What is the Internet address for the client host?
_____	What is an alias for the client host (optional)? Example: MYCLIENT
For each VMS User ID for which you wish a proxy on this client host...	
_____	What is a VMS user name to give the client host user proxy access? Example: SYSTEM
_____	What is a User ID (UID) on the client host? Example: 0 - must be an integer
_____	What is a Group ID (GID) on the client host? Example: 1 - must be an integer
NFS Disk Map Configuration	
For each disk to be served ... ¹	
_____	Disk to be served. Example: DUA100

¹These disks are assumed to be mounted. If not, the NFS configuration procedure can initialize and mount the disks the first time through. The next time the system is booted, all disks specified here must be mounted prior to UCX startup.

(continued on next page)

Table 7-1 (Cont.) Checklist for Configuring the NFS Server

Your Choice	Explanation
NFS Disk Map Configuration	
_____	<p>If you wish to restrict access to a host, what is its name?</p> <p>This must be a host defined in the Client Host and Proxy Configuration. Example: myclient.mytown.dec.com</p>
_____	<p>Do you wish to place a container file system on the disk, Yes or No</p>
_____	<p>If you do not wish to place a container file system on the disk, do you wish to enable name conversion, Yes or No</p> <p>Name conversion allows a client to create files with mixed case names and names containing characters which are invalid for OpenVMS file names.</p>

7.3 NFS Configuration Session

The following is an example of a configuration session.

```
Your StorageWorks Network Storage Array comes with Digital TCP/IP Services
for OpenVMS already installed. You may at this time allow the network array
to automatically activate licenses and configure this product for operation
as an NFS server. If you do opt for this automatic NFS configuration, you
will be asked a series of questions about your TCP/IP environment. Then
just those portions of the Digital TCP/IP Services for OpenVMS needed for
NFS and PCNFS operation will be configured. If you wish to use other
features of Digital TCP/IP Services for OpenVMS, you are free to use
UCX$CONFIG to do so later at your own convenience.
```

```
Do you wish to proceed with NFS configuration? [Y]:yes
```

```
TCP/IP Network Configuration Procedure
```

```
This procedure helps you define the parameters required
to run Digital TCP/IP Services for OpenVMS on this system.
```

```
The 'NFS configuration' option is selected.
```

```
The 'silent mode' ENABLE cluster wide option is also selected.
```

```
Checking TCP/IP Services for OpenVMS configuration database files.
```

```
Creating SYS$COMMON:[SYSEXE]UCX$SERVICE.DAT
```

```
Creating SYS$COMMON:[SYSEXE]UCX$HOST.DAT
```

```
Creating SYS$COMMON:[SYSEXE]UCX$NETWORK.DAT
```

```
Creating SYS$COMMON:[SYSEXE]UCX$ROUTE.DAT
```

```
For security purposes the Proxy Database is created with no
world access. This denies nonprivileged users the use of
commands which access the Proxy Database.
```

```
Creating SYS$COMMON:[SYSEXE]UCX$PROXY.DAT
```

```
Do you want to grant world read access to nonprivileged users? [N]no
```

```
Creating SYS$COMMON:[SYSEXE]UCX$CONFIGURATION.DAT
```

TIME ZONE Configuration

Please answer the following questions to configure your time zone.

Select your geographic area using the options in the table below.

Selection Geographic Area

```
-----  
u  USA  
ea  Eastern Australia  
ca  Central Australia  
wa  Western Australia  
ee  Eastern Europe  
ce  Central Europe  
we  Western Europe  
g   Greenwich  
n   none of the above  
-----
```

Enter your choice: u

Does your area alternate between Daylight Savings and Standard time [Y]?

Select the time zone for your area, using the options listed in the table below.

Selection Time Zone

```
-----  
e  Eastern  
c  Central  
m  Mountain  
p  Pacific  
g  Greenwich  
-----
```

Enter your choice: m

DOMAIN Configuration

Internet domain example:

In fully qualified host name 'myclient.mytown.dec.com',
the Internet domain is mytown.dec.com

Enter Internet domain: cxo.dec.com

INTERFACE Configuration

The FDDI device(s) on your system are: FWA0: * Not Configured *

Start of configuration questions for Internet interface WF0.
WF0 is the FDDI device FWA0:

Enter Internet address for THIS node, TOPCPU: 16.253.32.112
Enter Internet address for OTHER node, BOTCPU [16.253.32.111] :
Enter Internet network mask for TOPCPU and BOTCPU [255.0.0.0]:
Enter broadcast mask for TOPCPU and BOTCPU [16.255.255.255]:

The following parameters will be used to define the
Internet interface WF0:

```
This Host name:      TOPCPU.cxo.dec.com  
This Internet address: 16.253.32.112  
  
Other Host name:     BOTCPU.cxo.dec.com  
Other Internet address: 16.253.32.111  
  
Network mask:        255.0.0.0  
Broadcast mask:      16.255.255.255
```

* Is the above correct [YES] ?

End of configuration questions for Internet interface WF0

Because there are two CPU nodes in this server, an Internet VMScluster will be configured so that the two cluster members appear to remote network hosts as a single virtual host. This allows failover and load balancing.

This Internet VMScluster is identified by its cluster host name and address, which must be different from those you have just provided for the individual CPU nodes.

Start of cluster host questions for Internet interface WF0.
WF0 is the FDDI device FWA0:

Unqualified host name example:

In fully qualified host name 'myclient.mytown.dec.com',
the unqualified host name is myclient

Enter unqualified CLUSTER host name: CLHOST
Enter Internet address for falbal: 16.253.32.1125
Enter Internet network mask for falbal [255.0.0.0]:
Enter broadcast mask for falbal [16.255.255.255]:

The following parameters will be used to define the cluster host for Internet interface WF0:

CLUSTER host name: CLHOST.cxo.dec.com
Internet address: 16.253.32.115
Network mask: 255.0.0.0
Broadcast mask: 16.255.255.255

* Is the above correct [YES] ?

End of cluster host questions for Internet interface WF0

PORTMAPPER SERVER Configuration

Service is not defined in the SYSUAF.

Service is not defined in the UCX\$SERVICE database.

Creating UCX\$AUX identifier with a value of 375

%%%%%%%% OPCOM 27-MAR-1996 13:06:56.15 %%%%%%%%%

Message from user AUDIT\$SERVER on RLRBAL

Security alarm (SECURITY) and security audit (SECURITY) on RLRBAL, system id: 64
531

Auditable event: Identifier added
Event time: 27-MAR-1996 13:06:56.13
PID: 20200212
Process name: SYSTEM
Username: SYSTEM
Process owner: [SYSTEM]
Terminal name: OPA0:
Image name: DSA0:[SYS0.SYSCOMMON.][SYSEXE]AUTHORIZE.EXE
Identifier name: UCX\$AUX
Identifier value: %X00FDFFFF
Attributes: none

.
. .
. .
. .

Approximately how many NFS Clients will be served? [10]: 32

NFS SERVER Configuration

```

Service is not defined in the SYSUAF.
Service is not defined in the UCX$SERVICE database.
%%%%%%%%% OPCOM 27-MAR-1996 13:07:32.85 %%%%%%%%%%
Message from user AUDIT$SERVER on RLRBAL
Security alarm (SECURITY) and security audit (SECURITY) on RLRBAL, system id: 64
531
Auditable event:      System UAF record addition
Event time:          27-MAR-1996 13:07:32.83
PID:                 20200212
Process name:        SYSTEM
Username:            SYSTEM
Process owner:       [SYSTEM]
Terminal name:       OPA0:
Image name:          DSA0:[SYS0.SYSCOMMON.][SYSEXE]AUTHORIZE.EXE
Object class name:   FILE
Object name:         SYS$COMMON:[SYSEXE]SYSUAF.DAT;1
User record:         UCX$NFS
Account:             New:      UCX$NFS
                    Original:
Batch access:        New:      Primary: (no access), Secondary: (no access)
                    Original: Primary: (full access), Secondary: (full acc
ess)
DEFCLI:             New:      DCL
                    Original: DCL
CLITABLES:          New:      DCLTABLES
                    Original: DCLTABLES
Default Device:      New:      SYSS$SYSDEVICE:
                    Original: SYSS$SYSDEVICE:
Dialup access:       New:      Primary: (no access), Secondary: (no access)
                    Original: Primary: (full access), Secondary: (full acc
ess)
Default Directory:   New:      [UCX$NFS]
                    Original: [USER]
Flags:               New:      RESTRICTED
                    Original: DISUSER
Local access:        New:      Primary: (n
*TRUNCATED* -- If audits are enabled for this class, the full message
                    can be examined with ANALYZE/AUD
.
.
.
.
Do you want to grant world read access to non-privileged users? [N]

PCNFS SERVER Configuration

```

```

Service is not defined in the SYSUAF.
Service is not defined in the UCX$SERVICE database.
%%%%%%%%% OPCOM 27-MAR-1996 13:08:03.01 %%%%%%%%%%
Message from user AUDIT$SERVER on RLRBAL
Security alarm (SECURITY) and security audit (SECURITY) on RLRBAL, system id: 64
531
Auditable event:      System UAF record addition
Event time:          27-MAR-1996 13:08:02.99
PID:                 20200212
Process name:        SYSTEM
Username:            SYSTEM
Process owner:       [SYSTEM]
Terminal name:       OPA0:
Image name:          DSA0:[SYS0.SYSCOMMON.][SYSEXE]AUTHORIZE.EXE
Object class name:   FILE
Object name:         SYS$COMMON:[SYSEXE]SYSUAF.DAT;1
User record:         UCX$PCNFS
Account:             New:      UCX$PCNF
                    Original:
Batch access:        New:      Primary: (no access), Secondary: (no access)
                    Original: Primary: (full access), Secondary: (full acc
ess)
DEFCLI:             New:      DCL
                    Original: DCL
CLITABLES:          New:      DCLTABLES
                    Original: DCLTABLES
Default Device:      New:      SYSS$SYSDEVICE:
                    Original: SYSS$SYSDEVICE:
Dialup access:       New:      Primary: (no access), Secondary: (no access)
                    Original: Primary: (full access), Secondary: (full acc
ess)
Default Directory:   New:      [UCX$PCNFS]
                    Original: [USER]
Flags:              New:      RESTRICTED
                    Original: DISUSER
Local access:        New:      Primary:
*TRUNCATED* -- If audits are enabled for this class, the full message
                  can be examined with ANALYZE/AUD
.
.
.
.

NFS configuration has successfully completed.

```

7.4 Client Authorization and Served Disk Mapping Configuration

The following is an example of a Client Host and Proxy configuration session.

Client Host and Proxy Configuration

```

You will be prompted for a fully qualified host name, its
Internet address, and optionally an alias. Then for this
host, you can enter a series of UID/GID to VMS username proxies.
When you have completed entering this list of proxies for one
host, and you have terminated the list with a blank entry, you
will be prompted for another fully qualified host name -- and
so on, until you have identified all authorized client hosts.

```

```

Enter the fully qualified name of a client host: HOST.cxo.dec.com
Enter Internet address for HOST.cxo.dec.com: 16.82.16.215
Enter an alias for HOST.cxo.dec.com, or blank for none []: HOST

```

Client host HOST.cxo.dec.com has been added. Now you will be prompted for a series of VMS username to UID/GID mappings. This series can be terminated by a null entry. If a VMS username doesn't yet exist, the proxy will be entered, but not used, until the username is added with AUTHORIZE.

Enter a VMS username to give a HOST.cxo.dec.com user proxy access []: system
Enter a User ID (UID) on client HOST.cxo.dec.com []: 0
Enter a Group ID (GID) on client HOST.cxo.dec.com []: 1
Add proxy for uid/gid 0/1 from HOST.cxo.dec.com as user systest?

Proceed (Y), Skip (N), or make this proxy for all hosts (*) [Y]:
Enter a VMS username to give a HOST.cxo.dec.com user proxy access []: systest
Enter a User ID (UID) on client HOST.cxo.dec.com []: -2
Enter a Group ID (GID) on client HOST.cxo.dec.com []: -2

Add proxy for uid/gid -2/-2 from HOST.cxo.dec.com as user systest?

Proceed (Y), Skip (N), or make this proxy for all hosts (*) [Y]:
Enter a VMS username to give a HOST.cxo.dec.com user proxy access []:

End of proxy entry for client host HOST.cxo.dec.com.

Enter the fully qualified name of a client host: HOST1.cxo.dec.com
Enter Internet address for HOST1.cxo.dec.com: 16.63.32.110
Enter an alias for HOST1.cxo.dec.com, or blank for none []: HOST1

Client host HOST1.cxo.dec.com has been added. Now you will be prompted for a series of VMS username to UID/GID mappings. This series can be terminated by a null entry. If a VMS username doesn't yet exist, the proxy will be entered, but not used, until the username is added with AUTHORIZE.

Enter a VMS username to give a HOST1.cxo.dec.com user proxy access []: system
Enter a User ID (UID) on client HOST1.cxo.dec.com []: 0
Enter a Group ID (GID) on client HOST1.cxo.dec.com []: 1

Add proxy for uid/gid 0/1 from HOST1.cxo.dec.com as user system?

Proceed (Y), Skip (N), or make this proxy for all hosts (*) [Y]:
Enter a VMS username to give a HOST1.cxo.dec.com user proxy access []: system
Enter a User ID (UID) on client HOST1.cxo.dec.com []: -2
Enter a Group ID (GID) on client HOST1.cxo.dec.com []: -2

Add proxy for uid/gid -2/-2 from HOST1.cxo.dec.com as user system?

Proceed (Y), Skip (N), or make this proxy for all hosts (*) [Y]:
Enter a VMS username to give a HOST1.cxo.dec.com user proxy access []:

End of proxy entry for client host HOST1.cxo.dec.com.

Enter the fully qualified name of a client host:

End of client host entry.

The following is an example of the NFS Disk Map and Export session:

NFS Disk Map and Export Configuration

You will be prompted for a list of disks to be served to your previously identified client systems, and given the option to create UNIX container files on them. Each disk may then be exported for serving to one or all of these client systems. A pathname will be assigned, allowing you to NFS mount the served disk or container on your client(s). When you have completed designating disks, a null entry will complete this portion of the configuration.
Enter disk to be served []: dua450

Device Name	Device	Status	Error Count	Volume Label	Free Blocks	Trans Count	Mnt Cnt
\$13\$DUA450:	(NRCP1)	Online	0				

\$13\$dua450: is not mounted. OK to VMS INIT and MOUNT [Y]:
 %MOUNT-I-MOUNTED, U13DUA450: mounted on _\$13\$DUA450: (NRCP2)
 If you wish to restrict access to a host enter its name []:
 %NSA-I-ANYHOST, Any client host can access \$13\$dua450:.

Placing a container file system on a disk provides maximum UNIX functionality. Not doing so allows file sharing with VMS applications and better performance.

Do you wish to place a container file system on \$13\$dua450: [Y]: No

Name conversion is when a client can create files with mixed case names and names containing characters that are invalid for OpenVMS file names. The server converts such names to valid OpenVMS file names, and reverses the conversion when displaying the file names back to a non-OpenVMS client.

Do name conversion for \$13\$dua200: [Y]:

Enter disk to be served []:

End of disk serving configuration.

NFS configuration has successfully completed.

Do you wish to start NFS server now? [Y]: Yes

Integrating the Network Array into the Network System

After you have customized the OpenVMS Alpha operating system parameters and configured the storage system, you can complete the installation by connecting the network array to the network interconnect and by performing a few additional tasks. In addition, there are other optional postinstallation tasks that you may want to complete. This chapter discusses each task, required and optional, to complete the installation.

8.1 Connecting to the Network Interconnect

After the software customization and storage configuration have been completed, you can connect the StorageWorks Network Storage Array to the network interconnect and boot the network storage array so that it becomes a member of the network containing the client nodes. Perform the following steps to accomplish the task:

1. Shut down the network array OpenVMS Alpha operating system as follows:
 - a. Log in to a privileged account such as the system manager account established earlier.
 - b. Execute the system shutdown command file. For example, entering the following command invokes a system shutdown:

```
$ @SYS$SYSTEM:SHUTDOWN
```

Note

When issuing a SHUTDOWN command on one server processor, the other server processor may lock up. If the system locks for one minute, type INIT or BOOT on the console for the server processor which was just shut down.

- c. Answer the questions asked during the network array's execution of the shutdown command file and wait until the system has completed a logical shutdown.
2. Connect the network array to the network interconnect as follows:
 - a. Remove the dust covers attached to the network connections of the network array.
 - b. Place the dust covers in a protective container such as a resealable plastic bag. Place the container in a safe place for later use if you disconnect the network array from the network interconnect.
 - c. Connect the network to the network interconnect.

3. Cycle power for the server processor by pressing its power switch off and then on again. After the power-on self-tests complete and the system prompt >>> appears, type B to reboot the system. (See Example 3-1 for an example display of the power-on self-tests.)

As the network array boots, it will automatically request membership in the network containing the client nodes. On completion of the boot process, the network array will be able to serve its storage devices to the client nodes.

At this point, you have set the network array operating system parameters, configured your StorageWorks storage devices, and connected the network array to your network interconnect.

8.2 Accessing Important Documentation

Before completing the StorageWorks Network Storage Array installation procedure, you may want to access and print some documentation provided on the network array system disk. For example, if this is the first Alpha-based node joining your network, you should print the *OpenVMS Alpha Version 6.2 Upgrade and Installation Manual*, located on the network array's system disk. This manual provides detailed information on system halt, shut down, and boot procedures.

Some frequently used documentation for OpenVMS Alpha software is located on the StorageWorks Network Storage Array's system disk in the [DOCUMENTATION.V0xx] directory where xx represents the OpenVMS operating system software version number. For example, for Version 6.2, the directory is [DOCUMENTATION.V062]. Documentation on the operating system CD-ROM is in directory [ALPHA062.DOCUMENTATION].

When needed, take these steps to print the documentation:

1. Copy the appropriate document from SYSSYSDEVICE (the network array's system disk) or from the documentation CD-ROM to one of the devices on the server that can be accessed by another node.
2. From another node, print the file on an available printer.

8.3 Ensuring Viability of the Write-Back Cache

The write-back cache module installed in your network array contains batteries that were completely charged at the factory. It is normal for these batteries to discharge slightly in shipment.

The network array's write-back cache and RAID features require fully charged batteries to maintain absolute data integrity. After installation, these advanced features may not be available until the batteries have had an opportunity to completely recharge. The charging process may take up to 4 hours to complete.

8.4 Ensuring VMScluster System Compatibility

Note

The remainder of this chapter refers only to StorageWorks Network Storage Arrays installed in a VMScluster environment.

After installing the StorageWorks Network Storage Array software, there remain some tasks that will ensure the system operates correctly in the cluster. The required tasks are as follows:

- Install the Volume Shadowing Interoperability (I14Y) Kit on VAX and/or Alpha nodes in your cluster as appropriate. See the release notes (*ALPSHADnn_062.RELEASE_NOTES* and *VAXSHADnn_062.RELEASE_NOTES* where *nn* refers to the kit version) found on the network array's system disk for reference information. Refer to Section 4.2.2 for additional information on volume shadowing interoperability.
- Update other nodes in the cluster as required by any changes to the quorum scheme caused by installation of the StorageWorks Network Storage Array. You must update the EXPECTED_VOTES value in each cluster member's MODPARAMS.DAT file, then run AUTOGEN on each node followed by shutting down and rebooting the cluster. For example, if you add an HS121 and a quorum disk, then each cluster member will need to have its EXPECTED_VOTES parameter increased by 3. See *VMScCluster Systems for OpenVMS* for additional information on quorum issues.
- Run AUTOGEN on the network array to update cluster parameters. Refer to Section 9.5 for details on running AUTOGEN.

8.5 Activating Shadowing Functionality on the Network Array

By default, shadowing functionality is enabled on the network array. For additional information about shadowing system parameters and managing shadowing in a VMScCluster environment, refer to the manual *Volume Shadowing for OpenVMS*.

8.6 Installing and Registering Licenses

The StorageWorks Network Storage Array ships with licenses included as paper Product Authorization Keys (PAK) for the OpenVMS Alpha operating system (OPENVMS-ALPHA), volume shadowing (VOLSHAD), OpenVMS Clustering (VMSCCLUSTER), and DEC TCP/IP Services (UCX).

Note

If available, license PAKs will automatically be installed by the Software Customization Procedure. If license PAKs were not available from manufacturing, complete the following procedure to install and register your licenses.

You must perform the following tasks to use the software products included with the network array (including the OpenVMS operating system):

1. Set up the license database for each server node.
2. Register each of the software product licenses provided with the Storage Array in the cluster license database.

Note

It is important to note that in a VMScluster, one of the following conditions MUST be maintained:

- There is a single, common license database.
 - If there are multiple databases, then all must be identical.
-

8.6.1 Setting Up the License Database on a Cluster

The default license database for an unmodified OpenVMS operating system is on the system disk in SYSSCOMMON:[SYSEXEC]LMF\$LICENSE.LDB. If you move the database, you must define the logical name LMF\$LICENSE at the system level to point to the new database. Place permanent systemwide logical name definitions in the file SYSSCOMMON:[SYSMGR]SYLOGICALS.COM (see Section 8.6.1.1).

In a VMScluster environment, each node in the cluster must access either the same license database or an identical copy of the database seen by the other nodes.

If there is a disk in the cluster that will always be available to the cluster nodes, then it is normally convenient to put your common license database on that disk (which may be the system disk for some of the cluster nodes). For any systems that boot from a separate system disk, you must redirect LMF to the license database. This is done by defining the logical name LMF\$LICENSE to point to the common license database.

If it is not possible to provide a common disk for all cluster nodes, then you must keep separate identical license databases. Whenever one database is modified, you must copy it to the other locations in order to update the other databases.

Further information about managing software licenses on OpenVMS VAX and OpenVMS Alpha systems can be found in the system management manual *OpenVMS License Management Utility Manual*.

8.6.1.1 Common License Database

After the first reboot of the network array server processor during the Software Customization Procedure (while the network array is not yet part of the cluster), modify SYSSMANAGER:SYLOGICALS.COM to do the following:

1. Mount the disk where the license database is located.
2. Define the logical LMF\$LICENSE to point to the database.

Example:

In SYSSCOMMON:[SYSMGR]SYLOGICALS.COM, add the following lines:

```
$ MOUNT/SYSTEM/NOASSIST $17$DKA23: NSTAR 1$SYSTEM
$ DEFINE/SYSTEM/EXEC LMF$LICENSE-
  $17$DKA23: [VMS$COMMON.SYSEXEC]LMF$LICENSE.LDB
```

8.6.1.2 Separate License Databases

Obtain a copy of one of the cluster license databases. This step should be done after the second reboot of each network array server processor during the Software Customization Procedure so that the network array will be part of the existing cluster and will be able to mount disks where a copy of the database is located.

Example:

After rebooting the server processor to bring the network array into the cluster, log into the network array's operating system through the console port and enter the following commands:

```
$ MOUNT/NOASSIST $17$DKA23:NSTAR1$SYSTEM
$ COPY $17$DKA23:[VMS$COMMON.SYSEXE]LMF$LICENSE.LDB-
  SYS$COMMON:[SYSEXE]LMF$LICENSE.LDB
```

8.6.2 Registering the Licenses

After the second server processor reboot (the network array now being part of cluster, but having no licenses loaded):

1. Log into the system account on a console attached to the network array's server processor.
2. Register each license (three for each server node in a dual processor server) using the licensing procedure `SYSS$UPDATE:VMSLICENSE.COM`.

Example:

Assume in a dual-processor system that the two server nodes have SCS node names of `NSTAR1` and `NSTAR2` (selected during the Software Customization Procedure and written to the `SYSGEN` parameter `SCSNODE` on each node). In addition, assume the following authorization numbers for the two `OPENVMS-ALPHA` PAKs:

```
ALS-WM-95099-5555
ALS-WM-95099-7777
```

Enter the following command:

```
$ @SYSS$UPDATE:VMSLICENSE
```

A menu of options will be displayed. Select option 1, "REGISTER a Product Authorization Key," and follow the instructions to register each license.

3. Modify each `OPENVMS-ALPHA` license to include the name of the specific server node on which this license is being installed.

Because licenses for the OpenVMS operating system software specify the `NO_SHARE` option on their PAKs, in a `VMScluster` environment, you must restrict each OpenVMS license to a single node. Use the license command `LICENSE MODIFY/INCLUDE=node_name`, specifying only one SCS node name for each OpenVMS license.

Example:

```
$ LICENSE MODIFY OPENVMS-ALPHA-
  /INCLUDE=NSTAR1/AUTHORIZATION=ALS-WM-95099-5555
$ LICENSE MODIFY OPENVMS-ALPHA-
  /INCLUDE=NSTAR2/AUTHORIZATION=ALS-WM-95099-7777
```

4. Load the licenses by either rebooting the server processor at this point or by using the following command to load the OPENVMS-ALPHA license:

```
§ LICENSE LOAD OPENVMS-ALPHA
```

If the site uses separate license databases, then the updated license database on each node in the network array must be copied to each place where a separate license database exists.

If the StorageWorks Network Storage Array is being added to an existing cluster, then the licenses may be registered all at one time in the existing cluster license database (or databases) before the network array is installed. The remaining task after the servers are installed is to ensure that each server node has access to the common database. This can be accomplished either by defining LMF\$LICENSE in SYLOGICAL.COM (as described above) or by copying the updated LMF\$LICENSE.LDB to SYSSCOMMON:[SYSEXE]LMF\$LICENSE.LDB on each server node. During subsequent reboots, the licenses will be loaded automatically on each server.

Supporting and Operating the StorageWorks Network Storage Array

This chapter discusses system management issues for the StorageWorks Network Storage Array such as preparing for and installing system management software, setting certain operational parameters, and backing up the network array's system disk.

Note

For optimal performance of network arrays, you should limit the applications you install on the server processors to storage management and system management applications. By installing and executing production software on the processors, you could degrade the performance of the network array.

You also should consider the kinds of operations you plan to perform with the network arrays. For example, to realize optimal performance of your network array, you should perform backup operations of online storage to devices controlled by the same network array.

9.1 Customizing the System

Manage the StorageWorks Network Storage Array processors as you would other VMScluster nodes with respect to maintaining a functional, highly reliable computing environment. You can customize system files and install DECnet software and layered products as needed.

Before you make changes to the files on the network array system disk, you should make a backup copy (see Section 9.3). If the changes you make cause configuration problems, you can restore your configuration to a known state.

9.1.1 Customizing System Files

You can edit system files that allow you to tailor your system management environment. If you created a quorum disk on the network array, you need to add a MOUNT command in the system startup file to mount the quorum disk during network array startup.

When customizing your environment, you may decide to set system parameters in the following command files:

- SYCONFIG.COM
- SYLOGICALS.COM

- SYLOGIN.COM
- SYSTARTUP_VMS.COM

The release notes, release notes addendum, and *A Comparison of System Management on OpenVMS Alpha and OpenVMS VAX* contain detailed information for notes and restrictions that might be relevant to your customization plans. Additional information about maintaining the network array's performance can be found in the *Guide to OpenVMS Performance Management*.

Note

Immediately after changing your server processor system parameters, you should make a backup copy to protect your modifications.

9.1.2 Configuring and Starting DECnet for OpenVMS Alpha Software

If you plan to run DECnet™ for OpenVMS Alpha software, you must do the following:

1. After you have registered the license for the DECnet for OpenVMS Alpha software, execute the interactive command procedure SYSSMANAGER:NETCONFIG.COM to automatically configure your system for networking. See the *DECnet for OpenVMS Guide to Networking* for instructions on using NETCONFIG.COM.
2. After you start the queue manager (see the *OpenVMS System Manager's Manual*), edit the commands in SYSSCOMMON:[SYSMGR]SYSTARTUP_VMS.COM that pertain to networking so that the DECnet for OpenVMS software starts automatically when your system is booted. Choose one of the following commands to start the network and remove the comment delimiter (!) from that command:

```
$! IF F$SEARCH("SYSSSYSTEM:NETACP.EXE") .NES. "" THEN @SYSSMANAGER:STARTNET
$! IF F$SEARCH("SYSSSYSTEM:NETACP.EXE") .NES. "" THEN SUBMIT SYSSMANAGER:STARTNET.COM
```

Both of the previous commands perform the same task. However, the first command executes STARTNET.COM and delays further processing until the procedure is completed; the second submits STARTNET.COM to a batch queue and continues executing SYSTARTUP_VMS.COM.

Note

Immediately after changing your server processor system parameters, you should make a backup copy to protect your modifications.

9.1.3 Installing Layered Products

The StorageWorks Network Storage Array ships from the factory with all the layered products required for its operation. If you decide to install additional layered products such as Storage Library System for OpenVMS or Host Based RAID, note the following:

- You must register the licenses for OpenVMS Alpha layered products such as Host Based RAID for OpenVMS Alpha software.

You can invoke the OpenVMS License utility by entering the following command:

```
$ @SYS$UPDATE:VMSLICENSE
```

(You also can use the LICENSE REGISTER command.)

For additional information about registering licenses, see the *OpenVMS License Management Utility Manual*.

- For general information about installing layered products, see the *OpenVMS System Manager's Manual*. For more specific information, see the documentation you received with each layered product.

The DECwindows components provided with the OpenVMS Alpha Version 6.2 operating system supply only DECwindows base support and workstation support files. To get full DECwindows support, you also must install the separate DECwindows™ Motif® for OpenVMS Alpha layered product, which supports both the Motif and XUI environments.

See the most recent version of the *DECwindows Motif for OpenVMS Installation Guide* for information about installing the separate DECwindows Motif for OpenVMS Alpha layered product. After the installation, follow the directions in that manual and in *Managing DECwindows Motif for OpenVMS Systems* to customize your DECwindows environment.

Note

Immediately after changing your server processor system parameters, you should make a backup copy to protect your modifications.

9.2 Booting a Server Processor

At various times, you may have to boot the server processor. You can boot a server processor from either the network array's system disk or the CD-ROM device.

After the processor has been halted or after it has been powered up, the terminal connected to the processor displays the >>> prompt. To determine the name of the device from which you want to boot the server processor, then boot the device, follow this procedure:

1. Determine the names of the devices available to the server processor. Enter the following command at the >>> prompt:

```
>>> SHOW DEVICE
```

A display similar to the following appears:

dka400.4.0.6.0	DKA400 ❶	RRD45 1645	(1)
dkb0.0.0.11.0	DKB0 ❷	RZ28 442D	(2)
dkb200.2.0.11.0	DKB200 ❸	RZ26L 442D	(3)
dkc100.1.0.12.0	DKC100	RZ28 442D	(2)
dva0.0.0.0.1	DVA0		
fra0.0.0.1002.0	FRA0	08-00-2B-A6-0D-73	
pka0.6.0.6.0	PKA0	SCSI Bus ID 6	
pkb0.6.0.11.0	PKB0	SCSI Bus ID 6	
pkc0.6.0.12.0	PKC0	SCSI Bus ID 6	

Where:

- ❶ DKA400 is the name of the server processor CD-ROM
 - ❷ By default, the server processor shadowed system disk is composed of DKB0 and DKC100
 - ❸ By default, the server processor quorum disk is DKB200
2. To boot the server processor, specify the device by name in the boot command. For example, the following command shows how to boot from the CD-ROM shown in the previous example:

```
>>> BOOT DKA400
```

System roots [SYS0] and [SYS1] are reserved for use by the two server processors in the Network Storage Array. Console parameters should be set on the two server processors as follows:

Top server processor:

```
BOOT_OSFLAGS = 1,0  
BOOTDEF_DEV = DKB0,DKC100
```

Bottom server processor:

```
BOOT_OSFLAGS = 0,0  
BOOTDEF_DEV = DKB0,DKC100
```

With these settings, to boot from the shadowed system disk:

```
>>> BOOT
```

To do a conversational boot on the top server processor:

```
>>> BOOT -FLAGS 1,1
```

To do a conversational boot on the bottom server processor:

```
>>> BOOT -FLAGS 0,1
```

9.3 Backing Up the Network Array System Disk

After you create a backup copy of the files on the system disk, you should keep the copy in a safe location. For complete information about backup operations, see the *OpenVMS Alpha Upgrade and Installation Manual*.

To ensure complete integrity of the backup of the network array shadowed system disk, complete the following procedure:

Note

Some system disk backups can tolerate open files (for example, operator log files, accounting log files, error log files, and so forth) that may not be consistent or up-to-date on the backup save set. If your system disk backups can tolerate open files, alternate procedures for backing up the array shadowed system disk without shutting down both server processors are discussed in *VMScLuster Systems for OpenVMS* and *Volume Shadowing for OpenVMS*.

- Examine the current state of the array shadowed system disk by entering:

```
$ SHOW DEVICE SYS$SYSDEVICE
```

You will see a report similar to the following:

Device Name	Device Status	Error Count	Volume Label	Free Blocks	Trans Count	Mnt Cnt
DSA100:	Mounted	0	SYS\$NSTAR1	2690884	190	2
\$13\$DKB0:	(NSTAR1) ShadowSetMember	0	(member of DSA100:)			
\$13\$DKC100:	(NSTAR1) ShadowSetMember	0	(member of DSA100:)			

- Shut down both server processors.

Note

Before shutting down the server processors, quiesce any client applications that are using storage on this network array. This includes both NFS and OpenVMS clients.

Note

If the client system is booted from a system disk served by the network array, shut down those systems before shutting down the network array, because the system disk will not be available to clients during the backup.

- Boot one of the array processors from the CD-ROM (see Section 9.2). This starts the menu-driven command procedure shown in the following example:

This example uses the shadowed system disk displayed by the SHOW DEVICE SYSSYSDEVICE command.

```
$ MOUNT /OVERRIDE=IDENTIFICATION DSA100: /SHADOW=( $\$13\$DKB0$ :, $\$13\$DKC100$ :)
$ MOUNT /FOREIGN target_drive
$ BACKUP /IMAGE /VERIFY DSA100: target_drive
```

In this example the system disk and a target tape device are mounted so the BACKUP command can create a backup tape.

```
$ MOUNT/OVERRIDE=IDENTIFICATION DKA200
$ MOUNT/FOREIGN MKA300
$ BACKUP/IMAGE/VERIFY DKA200: MKA300:APR_06_BACKUP.BCK/SAVE_SET
```

The BACKUP command creates a system disk that includes a set of volume parameters provided by Digital, including a CLUSTER_SIZE (disk access scheme) that is appropriate for your system. (The CLUSTER_SIZE refers to the way files are stored on the disk, not to VMScluster environments.) You can change most volume parameters later with the SET VOLUME command.

However, to change the CLUSTER_SIZE, you must back up the system disk to a disk that has been previously initialized with the CLUSTER_SIZE that you want. For more information about initializing a disk and using the BACKUP command, see the *OpenVMS System Manager's Manual* and the *OpenVMS System Management Utilities Reference Manual*, and see the description of the INITIALIZE and BACKUP commands in the OpenVMS DCL Dictionary.

After you complete the backup operation, do the following:

- Enter the LOGOUT command to exit from the DCL environment and return to the menu.
- Choose the shutdown option (4).
- After the shutdown completes, boot both server processors from the system disk.

9.4 Defining Packet Sizes

The StorageWorks Network Storage Array operating system is OpenVMS Alpha operating system Version 6.2. The SYSGEN parameter governing packet size has been preset to 4468. All other nodes in the same cluster and on the network interconnect should be set to the same packet size according to the following rules:

If the Operating System of the Node Is . . .	Then Set the SYSGEN Parameter . . .
OpenVMS VAX Version 6.0 or higher	NISCS_MAX_PKTSZ to 4468
OpenVMS Alpha Version 6.1 or higher	NISCS_MAX_PKTSZ to 4468
OpenVMS VAX Version 5.5-2	LRPSIZE to 4474

More information on packet sizes and related configuration considerations can be found in the manual titled *VMScluster Systems for OpenVMS*.

9.5 Running AUTOGEN

When you installed the operating system, the system executed the AUTOGEN.COM procedure to set the values of system parameters and the sizes of the page, swap, and dump files according to the system configuration. As a postinstallation procedure, you need to run the AUTOGEN.COM procedure again to properly tune the system.

Run AUTOGEN as follows:

1. After 24 hours of operation, run AUTOGEN in feedback mode and reboot the system.
2. Run AUTOGEN again in feedback mode two workdays later, and then reboot the system.
3. Digital recommends that you run AUTOGEN from SAVPARAMS through TESTFILES on a weekly basis thereafter, and examine AGEN\$PARAMS.REPORT to determine the need for additional changes.

Based on your examination of AGEN\$PARAMS.REPORT, you might need to modify parameter values in MODPARAMS.DAT. If so, note the following:

- Hardcoded values in MODPARAMS.DAT should not hinder AUTOGEN's ability to calculate feedback parameters. AUTOGEN generally does not reduce the value of parameters that allocate resources; it considers current parameter values to be minimum values, which means that you do not have to add MIN_* symbols to MODPARAMS.DAT.
- AUTOGEN does increase parameter values according to its calculations unless you have specified explicit or maximum values (by adding MAX_* symbols) in MODPARAMS.DAT.

For more information about the MODPARAMS.DAT file and about using AUTOGEN in general, see the *OpenVMS System Manager's Manual*. Also, see the *OpenVMS System Manager's Manual: Tuning, Monitoring and Complex Systems* for information on running AUTOGEN and for additional information on system tuning.

9.6 Balancing I/O Load on the Dual-Processor Network Arrays

When a device is accessed on a network array from multiple client nodes, performance can be enhanced if the access pathways used in device mounting are different. In the dual-processor network array model, the two processors and two or more HS1CP controllers typically offer four or more possible paths to a given disk device. By distributing mount connections to a device, bandwidth contention along a given path is minimized. When the load is balanced across the two server processors, the network array provides better performance than if one server processor was heavily loaded while the other was mostly idle.

Load balancing can be implemented in one or both of two domains: at the device bus level, and at the server/network (mount path) level.

9.6.1 Device-Level Balancing

To ensure that all resources connected to the device bus are utilized fairly equally in network array's with dual-redundant device channel processors, the HS1CP provides a means of tuning their usage via the PREFER qualifier. This can be done on one of the HS1CP device channel processors:

```
HS1CP> SET unit PREFERRED_PATH=HS1CP_NAME
```

where *unit* specifies a device and *HS1CP_NAME* is THIS (for the HS1CP currently used) or OTHER (for the alternate HS1CP).

9.6.2 Mount Path-Level Balancing via the DCL PREFER Command

A DCL-level command, PREFER, can be incorporated into the DCL command table to provide a means of explicitly specifying a preferred disk path at mount time, or to change an existing mount path. You can use the PREFER command to set up device paths for each VMScluster node to take maximum advantage of the pathways available.

The PREFER command files should be available in the SYS\$EXAMPLES directory of your system disk (as PREFER.*). To set up the PREFER command:

1. Place PREFER.EXE in the SYSS\$SYSTEM directory. (Another directory may be used; edit the 'image' line in file PREFER.CLD to specify the directory.)
2. To define PREFER in the local process DCL command table, issue the command:

```
$ SET COMMAND PREFER
```

To define PREFER in the systemwide DCL command table, issue the following command from an account with CMKRNL privilege:

```
$ SET COMMAND /TABLE=SYS$LIBRARY:DCLTABLES -  
/OUTPUT=SYS$COMMON:[SYSLIB]DCLTABLES PREFER
```

VMS Install is needed to make PREFER available to other users, again requiring CMKRNL privilege:

```
$ INSTALL  
INSTALL> REPLACE SYS$LIBRARY:DCLTABLES  
INSTALL> EXIT
```

Note

The Install utility commands must be executed from all nodes in a VMScluster system. If your system configuration deviates from the standard usage of DCLTABLES, modify the above commands to suit your site. For further information on defining command verbs or using the Install utility, consult the VMS documentation set.

To use the preferred path functionality, enter the command:

```
$ PREFER unit/HOST=host[/FORCE]
```

where *unit* specifies a device and *host* is the name of the server node that constitutes the preferred path. /FORCE is an optional qualifier described in following paragraphs.

In the following example, device \$10\$DUA10: has a primary path through node NSTAR1 and a secondary path through node NSTAR2. To select NSTAR1 as the primary path, enter the following command:

```
$ PREFER $10$DUA10:/HOST=NSTAR1
```

This command sets the preferred path on the local node so that the next mount will use the selected path. If the /CLUSTER qualifier is used on the MOUNT command, all nodes in the cluster will use the selected path. Note that no change is made to the SHOW DEVICE output.

The /FORCE qualifier is used to select a preferred path for mounted disks. If the disk is mounted (not /FOREIGN), then the /FORCE qualifier will force the drive into mount verification and move it to the new controller. For example:

```
$ PREFER $10$DUA10:/HOST=NSTAR1/FORCE
```

The path used to remount the disk will be the preferred path of the node that performs the mount verification. No other nodes in the VMScluster system will alter their paths to the device. For proper operation, the /FORCE qualifier should be issued only from a host on which the device is mounted.

To select a preferred path for all nodes in a VMScluster system, use the VMS SYSMAN Utility to set the preferred path on all nodes in the cluster before executing the PREFER command with the /FORCE qualifier.

Note

The PREFER command must be defined as a DCL command verb on all nodes in the VMScluster system before using it within SYSMAN.

In the following example, the SYSMAN PREFER command sets the preferred path on all nodes so that the next mount command from any node will use the selected path. The DCL PREFER command using the /FORCE qualifier causes the device to enter mount verification. After mount verification completes, the device will be remounted by all nodes through the preferred path, NSTAR1.

```
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO PREFER $10$DUA10:/HOST=NSTAR1
SYSMAN> EXIT
$ PREFER $10$DUA10:/HOST=NSTAR1/FORCE
```


You can install the OpenVMS Alpha Operating System, migrate to a shadowed system disk, or execute DCL commands and procedures to perform "standalone" tasks, such as backing up the system disk.

To get help at any time, type a question mark (?) and press Return.

Please choose one of the following:

- (1) Install the OpenVMS Alpha Operating System, Version V6.2-1H3
- (2) Migrate to shadowed system disk
- (3) Execute DCL commands and procedures
- (4) Shut down this system

Enter your choice: (1/2/3/4) 3

3. Enter 3, then press the Return key to begin a limited DCL environment that is identified by the DCL prompt (\$).
4. At the prompt (\$), enter the SHOW DEVICES command.
5. Examine the list of devices so you can determine which device is the source drive (the drive holding the backed up files you want to restore) and which device is your target drive (the disk on which you want the files restored).
6. When you have determined which devices will be the source drive and target drive, mount those devices (and any other output devices you plan to use) before you perform any restore operations. Enter the MOUNT commands in the following format:

```
$ MOUNT/OVERRIDE=IDENTIFICATION source-drive  
$ MOUNT/FOREIGN target-drive
```

where:

source-drive is the device holding the files you want to restore.
target-drive is the destination.

(Note, however, that you must use the MOUNT/FOREIGN command if the source drive is a tape device.)

7. Enter the BACKUP command in the following format:

```
$ BACKUP/IMAGE/VERIFY source-drive: target-drive:
```

(You also must include the save set name and the /SAVE_SET qualifier if the source drive is a tape device.)

In this example, a backup disk and a target disk are mounted so the BACKUP command can restore the system disk from the backup disk:

```
$ MOUNT/OVERRIDE=IDENTIFICATION DKA300  
$ MOUNT/FOREIGN DKA200  
$ BACKUP/IMAGE/VERIFY DKA300: DKA200:
```

In this example, a backup tape and a target disk are mounted so the BACKUP command can restore the system disk from the backup tape:

```
$ MOUNT/FOREIGN MKA300  
$ MOUNT/FOREIGN DKA200  
$ BACKUP/IMAGE/VERIFY MKA300:APR_06_BACKUP.BCK/SAVE_SET DKA200:
```

After you complete the restore operation, do the following:

- Enter the LOGOUT command to exit from the DCL environment and return to the menu.

- Choose the shutdown option (4).
- After the shutdown completes, boot from the network array's system disk.

10.2 Rebuilding the System Disk from the Network Array CD-ROM

Perform the following steps to rebuild the StorageWorks Network Storage Array OpenVMS Alpha operating system:

1. Shut down the network array.

Note

Before shutting down the server processors, quiesce any client applications that are using storage on this network array. This includes both NFS and OpenVMS clients.

Note

Any node in the VMScluster system whose system disk is served through the network array should also be shut down. Shut down that system before shutting down the network array.

2. Boot the network array from the CD-ROM (see Section 9.2). This starts the menu-driven command procedure shown in the following example:

```

OpenVMS (TM) Alpha Operating System, Version V6.2-1H3

Copyright (c) 1996 Digital Equipment Corporation. All rights reserved.

+-----+
|                                     |
|               Welcome to the       |
|      StorageWorks (TM) Network Storage Array      |
|               Standalone Environment               |
|                                     |
+-----+

System initialization will take about one minute. Please wait...
This Network Storage Array has two server processors.

+-----+
|                                     |
|      StorageWorks Network Storage Array CD Menu      |
|                                     |
+-----+

You can install the OpenVMS Alpha Operating System, migrate to a
shadowed system disk, or execute DCL commands and procedures to
perform "standalone" tasks, such as backing up the system disk.
To get help at any time, type a question mark (?) and press Return.
Please choose one of the following:

    (1) Install the OpenVMS Alpha Operating System, Version V6.2-1H3
    (2) Migrate to shadowed system disk
    (3) Execute DCL commands and procedures
    (4) Shut down this system

Enter your choice: (1/2/3/4) 1

+-----+
|                                     |
|      Install the OpenVMS Alpha Operating System, Version V6.2-1H3      |
|                                     |
+-----+

The installation procedure will ask a series of questions.

```


You can install the OpenVMS Alpha Operating System, migrate to a shadowed system disk, or execute DCL commands and procedures to perform "standalone" tasks, such as backing up the system disk.

To get help at any time, type a question mark (?) and press Return.

Please choose one of the following:

- (1) Install the OpenVMS Alpha Operating System, Version V6.2-1H3
- (2) Migrate to shadowed system disk
- (3) Execute DCL commands and procedures
- (4) Shut down this system

Enter your choice: (1/2/3/4) 4

Shutting down the system. Please wait...

SYSTEM SHUTDOWN COMPLETE

4. The StorageWorks Network Storage Array Software Customization Procedure will run when the network array is rebooted.

For further information on the operating system installation refer to the *OpenVMS Alpha Upgrade and Installation Manual*.

Network Storage Array NFS Matrix

Product	Special Comments
SUN NFS	<p>Default MOUNT_NFS values allow file copy operations to timeout during a server failover resulting in an empty (size=0) file on the target disk. Incrementally increase the timeout and retry values during NFS mount by issuing a UNIX command on the SUN system similar to the following example:</p> <pre data-bbox="641 735 1169 766"># mount -t nfs -o timeo=144, retry=15 ...</pre> <p>The following line must be added as the first line of the file, /etc/init.d/nfs.server, on the SUN system:</p> <pre data-bbox="641 850 1274 882">/usr/sbin/ndd -set /dev/ip ip_path_mtu_discovery 0</pre>
NetManage ChameleonNFS	<p>NetManage ChameleonNFS software must be Version 5.01 or greater.</p> <p>On redundant models of the Network Storage Array, each individual server node and its cluster alias should appear to the NetManage Client on the available-to-mount window. If the cluster alias does not appear (a rare occurrence), simply click on one of the individual server node entries.</p> <p>PC/NFS will not start automatically on UCX. Execute the following VMS command on each server processor on the network array:</p> <pre data-bbox="641 1144 1128 1176">\$ @sys\$manager:nsa\$nfs_pcnfs_kick.com</pre> <p>When a Windows NT node accesses an HSM shelved Windows NT file, the Windows NT node may receive the following invalid error message:</p> <pre data-bbox="641 1260 1136 1312">The <served-disk-file> is not a windows held file - or the FILE is corrupted.</pre> <p>Allow enough time for the HSM UNSHELVE operation on the server processor to complete and reissue the command.</p>

Network Storage Array Interconnect Cables

Table B-1 provides a list of cables to connect the SWXNA-*xx* Network Storage Array to the network.

Table B-1 StorageWorks Network Storage Array Interconnect Cables

Adapter	SAS Interconnect		DAS Interconnect	
	Digital Part#	Connector Type	Digital Part#	Connector Type
DEFPA-DA FDDI (DAS)	BN34D-1	SC-MIC	BN34D-1	SC-MIC
	BN34D-3	SC-MIC	BN34D-3	SC-MIC
	BN34D-10†	SC-MIC	BN34D-10†	SC-MIC
			BN34B-1	SC-SC
			BN34B-3	SC-SC
			BN34B-10	SC-SC
			BN34B-20	SC-SC
			BN34B-30	SC-SC
			BN34A-1	SC-ST
			BN34A-3	SC-ST
			BN34A-10	SC-ST
			BN34A-20	SC-ST
			BN34A-30	SC-ST
DEFPA-MA CDDI (DAS)			BN25H-03	UTP/RJ45

†For connections greater than 10 meters, use Computer Crafts, Inc. cables:

F1511110-20 SC-MC

F1511110-30 SC-MC

Key for Connector Type

- The SWXNA-*xx* FDDI adapter uses an SC-type connector.
- Bridges and concentrators use MIC-type connectors.
- A SAS adapter cannot be functionally attached directly to an FDDI ring.
- The SWXNA-*xx* FDDI adapter connects to an FDDI ring through bridges and concentrators using an SC-MIC connection. This adapter also can attach directly to an FDDI ring using SC- or ST-type connectors.
- Two SWXNA-*xx* FDDI adapter cables are needed to connect the SWXNA-*xx* FDDI adapter to an FDDI ring in a DAS-mode. Use only one cable for a SAS-mode connection.

StorageWorks Glossary

ac distribution

The method of distributing ac power in a cabinet.

ac power supply

A power supply designed to produce dc power from an ac input.

adapter

A device that converts the protocol and hardware interface of one bus type into that of another without changing the functionality of the bus. *See* **signal converter**.

American National Standards Institute

See **ANSI**.

ANSI

American National Standards Institute. An organization that develops and publishes electronic and mechanical standards.

array controller

A hardware/software device that facilitates communications between a host and one or more devices organized in an array. HS-family controllers are examples of array controllers.

BA350–Mx controller shelf

The StorageWorks controller shelf used for HS-family controller modules, cache modules, and shelf power units.

BA350–Sx SBB shelf

A StorageWorks shelf used for only power units and SBBs.

cable distribution unit

See **CDU**.

carrier

A standard, StorageWorks shelf-compatible, plastic shell into which a device can be installed. Sometimes called SBB carrier.

CDU

Cable distribution unit. The AC power entry device for StorageWorks cabinets. The unit provides the connections necessary to distribute ac power to cabinet shelves and fans.

CLI

Command line interpreter. Operator command line interface for the HS1CP firmware.

command line interpreter

See **CLI**.

controller

A hardware/software device that facilitates communications between a host and one or more devices. A controller translates bus protocols and hardware interfaces and adds functionality to the host/device communications.

controller shelf

A StorageWorks shelf designed to contain controller and cache memory modules.

data center cabinet

A generic reference to the large cabinets, such as the SW800 series, in which StorageWorks components can be mounted.

Diagnostics and Utilities Protocol

See **DUP**.

dual cabinet power configuration

A cabinet ac power configuration in which two ac sources and two ac power supplies are used to provide redundant dc power to each of the cabinet's SBB shelves.

dual shelf power configuration

A cabinet ac power configuration in which one ac source and two ac power supplies are used to provide redundant dc power to each of the cabinet's SBB shelves.

dual-redundant configuration

A controller configuration consisting of a primary and backup controller in one controller shelf. If the primary controller fails, the backup controller assumes control over the failing controller's devices.

DUP

Diagnostic and Utility Protocol. Host application software that allows a host operator terminal to connect to the controller's command line interpreter. *See also* **virtual terminal**.

failover

Failover is the process that takes place when one controller in a dual-redundant configuration fails, and the other controller takes over the direction of the storage subsystem. The other controller continues to direct the storage subsystem until the failed controller becomes operational or is replaced.

fan

An airflow device mounted in a StorageWorks cabinet.

field replaceable unit

See **FRU**.

filler panel

A sheet metal or plastic panel used to cover unused mounting areas in StorageWorks cabinets and shelves.

FRU

field replaceable unit. A hardware component that can be replaced in the field.

full-height device

A single device that occupies an entire 5.25 inch SBB carrier. StorageWorks full-height devices have an order number suffix of “-VA.”

half-height device

A device that occupies half of a 5.25 inch SBB carrier. Two half-height devices can be mounted in a 5.25 inch SBB carrier. The first half-height device is normally mounted in the lower part of the carrier. The second device is normally mounted in the upper part of the carrier.

logical units

A group of devices addressable as a virtual unit.

maintenance terminal

The operator terminal used to identify an HS-family controller, to enable its host paths, to define its subsystem configuration, and to check its status. The HS-family maintenance terminal interface is designed to accept any terminal conforming to EIA-423. The maintenance terminal is only required to configure a storage subsystem and is not required for normal operations.

nonredundant

A configuration in which there is no backup hardware in place for the hardware that is present.

qualified device

A device that has been fully tested in all appropriate StorageWorks hardware and software configurations, and is in complete compliance with Digital and country-specific standards, for example, FCC and TÜV.

quiesce

To make a bus inactive or dormant. The operator must quiesce SCSI bus operations, for example, during a device warm swap.

RAID

Redundant array of independent disks. A set of storage techniques devised to increase the performance and availability of a storage subsystem.

Redundant Array of Independent Disks

See **RAID**.

SBB

StorageWorks building block. A device housed in a standard StorageWorks SBB carrier. An SBB has a standard physical and electrical interface that is compatible with those of StorageWorks shelves and enclosures.

SBB shelf

StorageWorks building block shelf. A StorageWorks shelf, such as the BA350–Sx, designed to house plug-in SBB modules.

shelf brackets

Sheet metal components designed to attach and position StorageWorks shelves in their associated enclosures.

signal converter

A device that converts the protocol and hardware interface of one bus type into that of another without changing the functionality of the bus. *See* **adapter**.

single cabinet power configuration

A cabinet ac power configuration in which only one ac source and one ac power supply is used to supply dc power to the cabinet's SBB shelves.

skirt

A trim panel designed to mount around the base of the cabinet.

storage sets

A grouping of disk drives that make up a new distinct container.

StorageWorks

Digital's family of modular data storage products that allows customers to design and configure their own storage subsystems. Components include power, packaging, cabling, devices, controllers, and software. Customers can integrate devices and array controllers in StorageWorks enclosures to form storage subsystems.

StorageWorks building block

See **SBB**.

stripesets

In a RAID configuration, a virtual disk drive with its physical data spread across multiple physical disks. Stripset configurations do not include a data recovery mechanism.

supported device

A device tested as functionally compatible with an approved StorageWorks hardware and software configuration.

virtual terminal

A software path from an operator terminal on the host to the controller's CLI interface. The path can be established via the host port on the controller (using DUP) or via the maintenance port through an intermediary host (VCS). A virtual terminal is also sometimes called a host console.

warm swap

A controller function that allows devices to be added, removed, or replaced while the subsystem remains operational. All activity on the device's SCSI bus must normally be halted for the duration of the warm swap operation.

write-back cache

A cache configuration that increases the performance of host write requests. When the host requests a write operation, the cache writes the host's data first to the cache memory, completing the host's request quickly. It performs the slower operation of flushing the data to the external storage device at a later time. The host sees the write operation as complete when the data has reached the cache.

A

AUTOGEN, 9-8

B

Backup

- network array system disk, 9-4
- system disk restore, 10-1

C

Cabinet

- applying power, 3-12
- inspecting, 3-9
- leveling, 3-4, 3-8
- placement of, 3-7
- ramp unloading procedures, 3-4
- site preparation, 2-1
- unloading, 3-4
- unpacking, 3-1

Cache Manager Screen

- using, 6-3

Cacher

- caching disks, 6-4
- resetting statistics, 6-5
- starting/stopping, 6-3
- statistics, 6-4
- summary, 6-4

CD-ROM

- contents, 1-15
- handling, 1-16
- system disk rebuild, 10-3
- using, 1-15

CFMENU, 5-2

- configuring storage devices, 5-5
- terminal setup, 5-3

Configuration

- considerations, 5-1
- DECnet software, 9-2
- Network Storage Array, 5-1

Custom

- system files, 9-1
- systems, 9-1

D

DECnet software, 9-2

Device channel processor

- description, 1-14
- I/O demands, 1-14

Device load balance, 9-9

Documentation, 1-16

E

Environmental stabilization, 3-1

G

Grounding equipment, 2-2

H

Host-based cache, 6-1

I

I/O load balance, 9-8

Inspection

- cabinet installation, 3-9
- ramps, 3-4
- shipped equipment, 3-1

Installation

- base trim kit, 3-8
- layered products, 9-3

Interconnect cables, B-1

L

Layered products

- installation of, 9-3

Leveling cabinet, 3-8

M

Mount balance, 9-9

N

- Network File System, 7-1
- Network Storage Array
 - major components, 1-9
 - system overview, 1-1
- Network Storage Array processor
 - CD-ROM, 1-11
 - description, 1-10
 - floppy drive, 1-13
 - specifications, 1-7
 - system disk, 1-11
- Network Storage Array upgrades, 1-6
- NFS, 7-1
 - Authorization Configuration, 7-11
 - configuration session, 7-7
 - Served Disk Mapping Configuration, 7-11
- NFS Matrix, A-1
- NFS Server
 - Checklist, 7-3
 - Configuring, 7-2

P

- Packet size, defining, 9-7
- PC-NFS Daemon, 7-2
- PerfectCache, 6-1
 - caching disks, 6-4
 - disabling, 6-5
 - enabling, 6-5
 - Menu, 6-2
 - resetting statistics, 6-5
 - starting/stopping, 6-3
 - startup state, 6-5
 - statistics, 6-4
 - summary, 6-4
- Placement of cabinet, 3-7
- Power
 - applying cabinet power, 3-12
- Power system survey, 2-2

Q

- Quorum scheme, 4-2

R

- RAID support, 1-15
- Ramps
 - inspection, 3-4
 - unloading, 3-4
- Read/write cache, 1-14
- Removing cabinet from pallet, 3-4
- Requirements
 - grounding, 2-2
 - power, 2-1
 - safety, 2-3

- Requirements (cont'd)
 - service, 2-3
- Restore system disk, 10-1

S

- Shadowed System Disk, 1-11
- Site preparation, 2-1
- Software Customization Procedure, 4-3
- Specifications
 - environmental stabilization, 3-1
 - server processor enclosure, 1-7
- Stabilization thermal, 3-1
- Storage device support, 1-15
- StorageWorks Network Storage Array terminal port
 - connecting to, 3-11
- SWXNA-Bx/Gx description, 1-1
- SWXNA-Ax/Fx conceptual model, 1-4
- SWXNA-Ax/Fx description, 1-3
- SWXNA-Bx/Gx conceptual model, 1-2
- SWXNA-Cx conceptual model, 1-2
- SWXNA-Cx description, 1-2
- SWXNA-DA, HA description, 1-5
- SWXNA-Ex/Jx conceptual model, 1-5
- SWXNA-Ex/Jx description, 1-4
- System disk
 - backup, 9-4
- System Disk, 1-11
- System disk restore, 10-1
- System overview, 1-1

T

- Terminal setup
 - CFMENU, 5-3
- Thermal stabilization, 3-2
- Trim kit installing, 3-8

U

- Unloading cabinet, 3-4
- Unpacking cabinet, 3-1
- Upgrades, 1-6

V

- VMScluster
 - considerations, 4-2
 - quorum scheme, 4-2
 - volume shadowing, 4-3