
OpenVMS System Manager's Manual, Volume 2: Tuning, Monitoring, and Complex Systems

Order Number: AA-PV5NG-TK

June 2002

This manual is part two of a task-oriented guide to managing a Compaq *OpenVMS* system.

Revision/Update Information: This manual supersedes the *OpenVMS System Manager's Manual, Volume 2: Tuning, Monitoring, and Complex Systems*, OpenVMS Version 7.3

Software Version: OpenVMS Alpha Version 7.3-1
OpenVMS VAX Version 7.3

Compaq Computer Corporation
Houston, Texas

© 2002 Compaq Information Technologies Group, L.P.

Compaq, the Compaq logo, Alpha, OpenVMS, Tru64, VAX, VMS, and the DIGITAL logo are trademarks of Compaq Information Technologies Group, L.P. in the U.S. and/or other countries.

Microsoft, MS-DOS, Visual C++, Windows, and Windows NT are trademarks of Microsoft Corporation in the U.S. and/or other countries.

Intel, Intel Inside, and Pentium are trademarks of Intel Corporation in the U.S. and/or other countries.

Motif, OSF/1, and UNIX are trademarks of The Open Group in the U.S. and/or other countries.

Java and all Java-based marks are trademarks or registered trademarks of Sun Microsystems, Inc., in the U.S. and other countries.

All other product names mentioned herein may be trademarks of their respective companies.

Confidential computer software. Valid license from Compaq required for possession, use, or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.

Compaq shall not be liable for technical or editorial errors or omissions contained herein. The information in this document is provided "as is" without warranty of any kind and is subject to change without notice. The warranties for Compaq products are set forth in the express limited warranty statements accompanying such products. Nothing herein should be construed as constituting an additional warranty.

ZK6017

The Compaq *OpenVMS* documentation set is available on CD-ROM.

This document was prepared using DECdocument, Version 3.3-1b.

Contents

Preface	xix
15 Managing System Parameters	
15.1 Understanding System Parameters	15-2
15.1.1 Default, Current, and Active Values	15-3
15.1.2 Pages and Pagelets	15-4
15.2 Recommended Method for Changing Parameter Values	15-4
15.3 Converting Your Customized Parameter Settings for Use with AUTOGEN	15-5
15.4 Understanding the AUTOGEN Command Procedure	15-8
15.4.1 AUTOGEN Feedback	15-10
15.4.2 Feedback Report (AGEN\$PARAMS.REPORT)	15-11
15.4.3 AUTOGEN Phases	15-16
15.4.4 AUTOGEN Parameter File (MODPARAMS.DAT)	15-17
15.5 Modifying System Parameters with AUTOGEN	15-18
15.5.1 Controlling AUTOGEN's Parameter Settings with MODPARAMS.DAT	15-18
15.5.1.1 Increasing a Value with the ADD_ Prefix	15-19
15.5.1.2 Specifying a Minimum Value with the MIN_ Prefix	15-20
15.5.1.3 Specifying a Maximum Value with the MAX_ Prefix	15-20
15.5.1.4 Specifying an Absolute Value	15-20
15.5.1.5 Defining the Number of VAXcluster Nodes (VAX Only)	15-21
15.5.1.6 Defining the Number of Ethernet Adapters (VAX Only)	15-21
15.5.1.7 Presetting Parameter Values Before Adding Memory (VAX Only)	15-21
15.5.1.8 Overriding Parameters Related to DECnet	15-21
15.5.1.9 Values Set for NPAGEDYN and NPAGEVIR	15-22
15.5.2 Specifying a Minimum Required Age for Feedback (VAX Only)	15-22
15.5.3 Including an External Parameter File in MODPARAMS.DAT	15-22
15.5.4 Turning Off Logging of DCL Statements	15-23
15.6 Automating AUTOGEN Reports	15-23
15.6.1 Changing Parameter Values After Reviewing AUTOGEN Reports ...	15-25
15.7 Managing System Parameters with SYSMAN	15-26
15.7.1 Understanding Parameter Values and SYSMAN	15-26
15.7.2 Showing Parameter Values with SYSMAN	15-27
15.7.3 Modifying a Parameter File with SYSMAN	15-28
15.7.4 Modifying Active Values with SYSMAN	15-29
15.8 Managing System Parameters with SYSGEN	15-30
15.8.1 Understanding Parameter Values and SYSGEN	15-31
15.8.2 Showing Parameter Values with SYSGEN	15-32
15.8.3 Modifying the System Parameter File with SYSGEN	15-33
15.8.4 Modifying Active Values with SYSGEN	15-34
15.8.5 Creating a New Parameter File with SYSGEN	15-35

15.9	Modifying System Parameters with a Conversational Boot	15-36
15.10	Tuning BAP System Parameters	15-37

16 Managing Page, Swap, and Dump Files

16.1	Understanding Dump Files	16-2
16.1.1	Using the Page File to Store System Crash Dumps	16-3
16.1.2	Understanding Types of System Dumps	16-4
16.2	Understanding Page and Swap Files	16-5
16.3	Displaying Information About Page and Swap Files	16-6
16.4	Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files	16-7
16.4.1	Calculating System Dump File Size	16-7
16.4.2	Calculating Error Log Dump File Size	16-8
16.4.3	Calculating Page File Size	16-9
16.4.3.1	Representing Page File Size	16-10
16.4.3.2	Monitoring Page File Usage	16-10
16.4.3.3	Limiting Page File Space	16-11
16.4.4	Calculating Swap File Size	16-11
16.4.4.1	Representing Swap File Size	16-11
16.4.4.2	Monitoring Swap File Usage	16-11
16.5	Minimizing System Dump File Size When Disk Space Is Insufficient	16-12
16.5.1	Understanding the Order of Information in a Selective System Dump	16-13
16.5.2	Fine-Tuning the Order That Processes Are Written in Selective System Dumps (Alpha Only)	16-14
16.6	Writing the System Dump File to the System Disk	16-14
16.6.1	System Dump to System Disk on Alpha	16-14
16.6.2	System Dump to System Disk on VAX	16-15
16.7	Writing the System Dump File to an Alternate Disk	16-15
16.7.1	DOSD Requirements on Alpha Systems	16-15
16.7.2	DOSD Requirements on VAX Systems	16-18
16.8	Using SDA to Analyze the Contents of a Crash Dump	16-19
16.9	Using SDA CLUE Commands to Analyze Crash Dump Files (Alpha Only)	16-19
16.9.1	Understanding CLUE (Alpha Only)	16-19
16.9.2	Displaying Data Using SDA CLUE Commands (Alpha Only)	16-20
16.9.3	Using SDA CLUE with Dump Off System Disk (Alpha Only)	16-20
16.10	Using CLUE to Obtain Historical Information About Crash Dumps (VAX Only)	16-21
16.10.1	Understanding CLUE (VAX Only)	16-21
16.10.2	Displaying Data Using CLUE (VAX Only)	16-21
16.11	Saving the Contents of the System Dump File After a System Failure	16-22
16.12	Copying System Dump Files to Tape or Disk	16-23
16.13	Freeing Dump Information from the Page File	16-24
16.13.1	Freeing Dump Information on VAX and Alpha Systems	16-24
16.13.2	Usage Notes for Freeing Information on VAX and Alpha Systems	16-25
16.14	Installing Page and Swap Files	16-25
16.14.1	Installing Interactively	16-26
16.14.2	Installing in SYPAGSWPFILES.COM	16-26
16.15	Removing Page, Swap, and Dump Files	16-28
16.16	Creating and Modifying Page, Swap, and Dump Files	16-28

16.16.1	Using AUTOGEN (Recommended Method)	16-28
16.16.1.1	Controlling the Location of System Page, Swap, and Dump Files	16-29
16.16.1.2	Controlling the Size of System Page, Swap, and Dump Files in MODPARAMS.DAT	16-29
16.16.2	Using SWAPFILES.COM	16-32
16.16.3	Using SYSGEN	16-33
16.17	Understanding Process Dumps	16-35
16.17.1	Understanding Privileged Users and Access to Process Dumps (Alpha Only)	16-36
16.17.2	Granting Access to Process Dumps (Alpha Only)	16-36
16.17.3	Restricting Access to Process Dumps (Alpha Only)	16-37

17 Performance Considerations

17.1	Understanding Performance Management	17-1
17.2	Knowing Your Work Load	17-2
17.3	Choosing a Work Load Management Strategy	17-3
17.4	Distributing the Work Load	17-3
17.5	Understanding System Tuning	17-4
17.6	Predicting When Tuning Is Required	17-5
17.7	Evaluating Tuning Success	17-6
17.8	Choosing Performance Options	17-6
17.9	Expanding System Libraries	17-8
17.9.1	Determining Disk Space Available to Expand Libraries	17-9
17.9.2	Using the Library Decompression Utility (LIBDECOMP.COM)	17-9
17.9.2.1	Libraries on which LIBDECOMP.COM Operates	17-9
17.9.2.2	Using LIBDECOMP.COM Interactively	17-11
17.9.2.2.1	Listing Libraries	17-12
17.9.2.2.2	Expanding (Decompressing) Libraries	17-13
17.9.2.2.3	Reducing (Compressing) Libraries	17-15
17.9.2.3	Using LIBDECOMP.COM in Batch Mode	17-16
17.9.3	Using the LIBRARY Command with the /DATA Qualifier	17-16
17.10	Using INSTALL to Install Known Images	17-17
17.10.1	Understanding Images and Known Images	17-17
17.10.2	Understanding Known File Entries	17-18
17.10.3	Understanding Attributes You Can Assign to Known Images	17-18
17.10.4	Determining Which Images to Install	17-19
17.10.5	Installing Images to Improve Image Activation Performance	17-20
17.10.6	Installing Images with Shared Address Data	17-20
17.10.6.1	System-Provided Images	17-21
17.10.6.2	Application Images	17-21
17.10.7	Installing Images to Conserve Physical Memory	17-21
17.10.8	Installing Images to Enhance Privileges of Images	17-22
17.10.8.1	Privileged Executable Images	17-22
17.10.8.2	Privileged Shareable Images	17-22
17.10.9	Activating Images in a Privileged Context	17-23
17.10.10	Specifying File Names in INSTALL	17-24
17.10.11	Installing Images with INSTALL	17-24
17.10.12	Displaying Known Images with INSTALL	17-25
17.10.13	Defining Logical Names for Shareable Image Files	17-25
17.10.14	Removing Known Images	17-26
17.11	Reserved Memory Registry	17-26

17.11.1	Using the Reserved Memory Registry	17-27
17.11.1.1	Reserved Memory Registry Data File	17-27
17.11.1.2	AUTOGEN	17-27
17.11.1.3	Adding Entries to the Reserved Memory Registry	17-27
17.11.2	Removing Entries from the Reserved Memory Registry	17-29
17.11.2.1	Allocating Reserved Memory	17-29
17.11.2.2	Freeing Reserved Memory	17-29
17.11.2.3	Displaying Reserved Memory	17-30
17.11.2.4	Using Reserved Memory	17-31
17.11.2.5	Returning Reserved Memory	17-31
17.11.3	Application Configuration	17-32

18 Managing File System Data Caches

18.1	Understanding Caching	18-1
18.2	Understanding File System Data Caches	18-2
18.3	Disabling Caching Clusterwide	18-3
18.4	Mounting a Volume With Caching Disabled	18-3
18.5	Managing XFC (Alpha Only)	18-3
18.5.1	Controlling the Size of the Cache	18-4
18.5.1.1	Controlling the Minimum Cache Size	18-4
18.5.1.2	Controlling the Maximum Cache Size	18-5
18.5.1.3	Enabling a Static Cache Size	18-6
18.5.2	Controlling the Maximum Cached I/O Size	18-6
18.5.3	Disabling Caching for a File	18-7
18.5.4	Disabling Read-Ahead Caching	18-8
18.5.5	Monitoring Performance	18-8
18.5.5.1	System-Wide Statistics	18-8
18.5.6	Using XFC in a Mixed Architecture OpenVMS Cluster	18-9
18.6	Managing the Virtual I/O Cache	18-10
18.6.1	Understanding How the Cache Works	18-10
18.6.2	Selecting VIOC on an Alpha System	18-11
18.6.3	Controlling the Size of the Cache	18-11
18.6.4	Displaying VIOC Statistics	18-12
18.6.5	Enabling VIOC	18-13
18.6.6	Determining If VIOC Is Enabled	18-13
18.6.7	Memory Allocation and VIOC	18-13
18.6.8	Adjusting VIOC Size	18-14
18.6.9	VIOC and OpenVMS Cluster Configurations	18-14

19 Testing the System with UETP

19.1	Overview	19-1
19.1.1	Understanding UETP	19-1
19.1.2	Summary of How to Use UETP	19-2
19.2	Preparing to Use UETP	19-4
19.2.1	Logging In	19-4
19.2.2	Using the SYSTEST Directories	19-4
19.3	Setting Up the Devices to Be Tested	19-5
19.3.1	Check Your Devices	19-5
19.3.2	System Disk Space Required	19-5
19.3.3	How UETP Works on Disks	19-5
19.3.4	Prepare Disk Drives	19-6
19.3.5	Magnetic Tape Drives	19-6

19.3.6	Tape Cartridge Drives	19-7
19.3.7	Compact Disc Drives	19-7
19.3.8	Optical Disk Drives	19-7
19.3.9	Terminals and Line Printers	19-8
19.3.10	Ethernet Adapters	19-8
19.3.11	DR11-W Data Interface (VAX Only)	19-8
19.3.12	DRV11-WA Data Interface (VAX Only)	19-9
19.3.13	DR750 or DR780 (DR32 Interface) (VAX Only)	19-9
19.3.14	Second LPA11-K Device	19-10
19.3.15	Devices That Are Not Tested	19-10
19.3.16	OpenVMS Cluster Testing	19-10
19.3.17	Testing a Small-Disk System	19-11
19.3.18	DECnet for OpenVMS Phase	19-12
19.3.19	Vector Processors and the VVIEF (VAX Only)	19-12
19.4	Starting UETP	19-13
19.4.1	Running a Subset of Phases	19-13
19.4.2	Single Run Versus Multiple Passes	19-14
19.4.3	Defining User Load for Load Test	19-14
19.4.4	Report Formats	19-14
19.4.4.1	Long Report Format	19-15
19.4.4.2	Short Report Format	19-15
19.5	Stopping a UETP Operation	19-15
19.5.1	Using Ctrl/Y	19-15
19.5.2	Using DCL Commands	19-16
19.5.3	Using Ctrl/C	19-16
19.6	Troubleshooting: An Overview	19-16
19.6.1	Error Logging and Diagnostics	19-16
19.6.2	Interpreting UETP Output	19-17
19.6.3	Displaying Information on Your Screen	19-18
19.6.4	Example Screen Display (VAX Only)	19-18
19.6.5	Example Screen Display (Alpha Only)	19-19
19.6.6	Defining a Remote Node for UETP Ethernet Testing	19-20
19.6.7	Log Files	19-20
19.7	Troubleshooting: Possible UETP Errors	19-21
19.7.1	Summary of Common Failures	19-21
19.7.2	Wrong Quotas, Privileges, or Account	19-22
19.7.3	UETINIT01 Failure	19-24
19.7.4	UETVECTOR Failure (VAX Only)	19-25
19.7.5	Device Allocated or in Use by Another Application	19-25
19.7.6	Insufficient Disk Space	19-25
19.7.7	Incorrect Setup of an OpenVMS Cluster System	19-26
19.7.8	Problems During the Load Test	19-27
19.7.9	DECnet for OpenVMS Error	19-29
19.7.10	Errors Logged but Not Displayed	19-29
19.7.11	No PCB or Swap Slots	19-29
19.7.12	No Keyboard Response or System Disk Activity	19-30
19.7.13	Lack of Default Access for the FAL Object	19-31
19.7.14	Bugchecks and Machine Checks	19-31
19.8	UETP Tests and Phases	19-31
19.8.1	Initialization Phase	19-31

19.8.2	Device Test Phase	19-32
19.8.2.1	How the Device Phase Works	19-32
19.8.2.2	Running a Single Device Test	19-32
19.8.2.3	Format of UETINIDEV.DAT	19-33
19.8.2.4	Running a Test in Loop Mode	19-33
19.8.2.5	Functions of Individual Device Tests	19-34
19.8.3	System Load Test Phase	19-35
19.8.4	DECnet for OpenVMS Test Phase	19-36
19.8.4.1	Environment	19-36
19.8.4.2	How the DECnet Phase Works	19-37
19.8.5	Cluster-Integration Test Phase	19-38

20 Getting Information About the System

20.1	Understanding System Log Files	20-1
20.2	Understanding Error Logging	20-2
20.3	Using the Error Formatter	20-3
20.3.1	Restarting the ERRFMT Process	20-3
20.3.2	Maintaining Error Log Files	20-4
20.3.3	Using ERRFMT to Send Mail	20-4
20.3.3.1	Enabling and Disabling ERRFMT to Send Mail	20-5
20.3.3.2	Sending Mail to Another User	20-5
20.4	Using ERROR LOG	20-5
20.4.1	Understanding ERROR LOG	20-5
20.4.2	Producing Error Log Reports	20-7
20.4.3	Producing a Full Error Log Report	20-7
20.4.4	Using Other Error Log Report Options	20-8
20.5	Using DECEvent	20-8
20.5.1	Understanding DECEvent	20-9
20.5.2	Invoking and Exiting DECEvent	20-10
20.5.3	Using DECEvent Qualifiers	20-10
20.5.4	Using Additional DECEvent Commands	20-12
20.5.5	Producing DECEvent Reports	20-12
20.5.5.1	Producing a Full Report	20-12
20.5.5.2	Producing a Brief Report	20-13
20.5.5.3	Producing a Terse Report	20-14
20.5.5.4	Producing a Summary Report	20-15
20.5.5.5	Producing a Fast Error (FSTERR) Report	20-16
20.5.6	DECEvent Restrictions	20-16
20.6	Setting Up, Maintaining, and Printing the Operator Log File	20-17
20.6.1	Understanding the Operator Log File	20-17
20.6.2	Understanding OPCOM Messages	20-18
20.6.2.1	Initialization Messages	20-18
20.6.2.2	Device Status Messages	20-18
20.6.2.3	Terminal Enable and Disable Messages	20-18
20.6.2.4	User Request and Operator Reply Messages	20-20
20.6.2.5	Volume Mount and Dismount Messages	20-21
20.6.2.6	System Parameter Messages	20-21
20.6.2.7	Security Alarm Messages	20-22
20.6.2.8	Contents of an Operator Log File	20-22
20.6.3	Setting Up the Operator Log File	20-23
20.6.3.1	Creating a New Version of the Operator Log File	20-24
20.6.3.2	Specifying Logical Names	20-24
20.6.4	Maintaining the Operator Log File	20-25

20.6.5	Printing the Operator Log File	20-25
20.7	Using Security Auditing	20-26
20.7.1	Understanding Security Auditing	20-27
20.7.1.1	Security Audit Log File	20-27
20.7.1.2	Audit Log Files in Mixed-Version Clusters	20-28
20.7.2	Displaying Security Auditing Information	20-28
20.7.3	Delaying Startup of Auditing	20-29
20.7.4	Enabling Security Auditing for Additional Classes	20-29
20.7.5	Disabling Security Auditing	20-30
20.7.6	Enabling a Terminal to Receive Alarm Messages	20-30
20.7.7	Generating Security Reports	20-31
20.7.8	Creating a New Version of the Security Audit Log File	20-31
20.7.8.1	Creating a New Clusterwide Version of the Log File	20-31
20.7.8.2	Creating a New Node-Specific Version of the Log File	20-31
20.8	Monitoring Operating System Performance	20-32
20.8.1	Understanding MONITOR	20-32
20.8.1.1	MONITOR Classes	20-32
20.8.1.2	Display Data	20-34
20.8.1.3	Output Types	20-34
20.8.2	Invoking MONITOR	20-35
20.8.3	Using Live Display Monitoring	20-35
20.8.4	Using Live Recording Monitoring	20-37
20.8.5	Using Concurrent Display and Recording Monitoring	20-37
20.8.6	Using Playback Monitoring	20-38
20.8.7	Using Remote Playback Monitoring	20-40
20.8.8	Rerecording Monitoring	20-40
20.8.9	Running MONITOR Continuously	20-41
20.8.9.1	Using the MONITOR.COM Procedure	20-41
20.8.9.2	Using the SUBMON.COM Procedure	20-43
20.8.9.3	Using the MONSUM.COM Procedure	20-43
20.8.10	Using Remote Monitoring	20-45

21 Tracking Resource Use

21.1	Understanding Accounting Files	21-1
21.2	Determining Which Resources Are Being Tracked	21-2
21.3	Controlling Which Resources Are Tracked	21-3
21.4	Starting Up a New Accounting File	21-3
21.5	Moving the Accounting File	21-3
21.6	Producing Reports of Resource Use	21-4
21.7	Setting Up Accounting Groups	21-5
21.8	Monitoring Disk Space	21-6

22 OpenVMS Cluster Considerations

22.1	Understanding OpenVMS Cluster Systems	22-1
22.1.1	Setting Up an OpenVMS Cluster Environment	22-2
22.1.2	Clusterwide System Management	22-3
22.2	Using Compaq DECams to Analyze Data	22-4
22.3	Using SHOW CLUSTER	22-4
22.3.1	Understanding SHOW CLUSTER	22-4

22.3.2	Beginning to Use SHOW CLUSTER Commands	22-6
22.3.2.1	Viewing Information That Is Off the Screen	22-6
22.3.2.2	Exiting from a Continuous Display	22-7
22.3.2.3	Using SHOW CLUSTER Qualifiers	22-8
22.3.3	Adding Information to a Report	22-8
22.3.4	Controlling the Display of Data	22-10
22.3.4.1	Entering Commands to Display Data	22-10
22.3.4.2	Removing Broadcast Messages	22-10
22.3.4.3	Refreshing the Screen	22-10
22.3.5	Formatting the Display of Data	22-11
22.3.5.1	Removing Information from a Report	22-11
22.3.5.2	Modifying Field and Screen Size	22-11
22.3.5.3	Moving a Report	22-12
22.3.5.4	Scrolling a Report	22-13
22.3.6	Creating a Startup Initialization File	22-14
22.3.7	Using Command Procedures Containing SHOW CLUSTER Commands	22-15
22.4	Understanding SYSMAN and OpenVMS Cluster Management	22-16
22.5	Using SYSMAN to Manage Security	22-16
22.5.1	Modifying the Group Number and Password	22-16
22.6	Using the SYSMAN Command DO to Manage an OpenVMS Cluster	22-17

23 Network Considerations

23.1	OpenVMS Networking Software Options	23-2
23.2	Choosing Compaq Networking Software	23-2
23.3	Understanding Compaq TCP/IP Services for OpenVMS	23-4
23.3.1	Support for OpenVMS Cluster Systems	23-5
23.3.2	TCP/IP Services Management Tools and Utilities	23-5
23.4	Preparing to Join a TCP/IP Network	23-6
23.5	Installing and Configuring TCP/IP Services	23-6
23.6	Starting and Stopping TCP/IP Services	23-7
23.7	TCP/IP Services Documentation	23-7
23.8	Understanding DECnet-Plus for OpenVMS Networking Software	23-8
23.8.1	DECnet-Plus Node Names	23-9
23.8.2	Support for OpenVMS Cluster Systems	23-10
23.8.3	DECnet-Plus Management Tools and Utilities	23-10
23.9	Preparing to Join a DECnet-Plus Network	23-11
23.10	Installing and Configuring DECnet-Plus	23-11
23.11	Using DECnet Over TCP/IP	23-12
23.12	Moving Your DECnet Phase IV Network to DECnet-Plus	23-13
23.13	Starting and Stopping DECnet-Plus	23-13
23.14	DECnet-Plus Documentation	23-13

24 Managing the Local Area Network (LAN) Software

24.1	Understanding Local Area Networks	24-1
24.1.1	LAN Characteristics	24-2
24.1.1.1	Ethernet LANs	24-3
24.1.1.2	FDDI LANs	24-3
24.1.1.3	Token Ring LANs (Alpha Only)	24-3
24.1.1.4	ATM LANs (Alpha Only)	24-4
24.1.2	LAN Addresses	24-4
24.2	Managing Local Area Networks	24-5

24.3	Understanding the LANACP LAN Server Process	24-6
24.3.1	Running the LANACP LAN Server Process	24-7
24.3.2	Stopping the LANACP LAN Server Process	24-7
24.4	Understanding the LANCP Utility	24-7
24.4.1	Invoking and Running LANCP	24-7
24.4.2	LANCP Commands	24-8
24.4.3	LANCP Miscellaneous Functions	24-10
24.5	Managing LAN Devices	24-10
24.5.1	Displaying System Devices	24-10
24.5.2	Displaying Device Characteristics	24-11
24.5.3	Setting Device Characteristics	24-13
24.6	Managing the LAN Device Databases	24-21
24.6.1	Displaying Devices in the LAN Device Databases	24-21
24.6.2	Entering Devices into the LAN Device Databases	24-22
24.6.3	Deleting Devices from the LAN Device Databases	24-22
24.7	Managing the LAN Node Databases	24-23
24.7.1	Displaying Nodes in the LAN Node Databases	24-23
24.7.2	Entering Nodes into the LAN Node Databases	24-23
24.7.3	Deleting Nodes from the LAN Node Databases	24-25
24.8	Understanding LAN MOP	24-25
24.8.1	Coexistence with DECnet MOP	24-25
24.8.2	Migrating from DECnet MOP to LAN MOP	24-26
24.8.3	Using CLUSTER_CONFIG_LAN.COM and LAN MOP	24-26
24.8.4	Sample Satellite Load	24-27
24.8.5	Cross-Architecture Booting	24-28
24.9	Managing the LAN MOP Downline Load Service	24-28
24.9.1	Enabling MOP Downline Load Service	24-28
24.9.2	Disabling MOP Downline Load Service	24-28
24.9.3	Displaying the Status and Counters Data	24-28
24.9.4	Displaying the Status and Counters Data for Individual Nodes	24-29
24.9.5	Clearing the Counters Data	24-30
24.9.6	OPCOM Messages	24-30
24.9.7	Load Trace Facility	24-30
24.9.8	MOP Console Carrier	24-31
24.9.9	MOP Trigger Boot	24-32

25 Managing InfoServer Systems

25.1	Understanding InfoServer Functions	25-1
25.1.1	Automatic Service Policies for Multiple Servers	25-4
25.1.2	High-Availability Feature to Reduce Service Interruptions	25-4
25.1.3	Support for X Terminal Clients	25-5
25.2	Understanding LASTport Protocols	25-5
25.2.1	LASTport Transport Protocol	25-5
25.2.2	LASTport/Disk Protocol	25-6
25.2.3	LASTport/Tape Protocol	25-6
25.3	Establishing a Server Management Session	25-6
25.3.1	Server Management Commands	25-7
25.4	Understanding InfoServer Client for OpenVMS Functions	25-9
25.5	Understanding LASTCP Utility Functions	25-9
25.5.1	Invoking and Exiting the LASTCP Utility	25-10
25.5.2	LASTCP Command Summary	25-10
25.5.3	Starting InfoServer Client for OpenVMS Software Automatically	25-11

25.5.4	InfoServer Client Can Fail to Start If DECnet Is Started or Stopped	25-12
25.5.5	Multiple Controllers Configured But Not All Attached to Media (Alpha Only)	25-12
25.5.6	Startup Restrictions: PATHWORKS and RSM	25-14
25.5.7	Startup Restrictions: SYSMAN	25-14
25.5.8	User Account Requirements	25-14
25.5.9	System Parameter MAXBUF Requirement	25-14
25.6	Understanding LADCP Utility Functions	25-14
25.6.1	Invoking and Exiting the LADCP Utility	25-15
25.6.2	LADCP Command Summary	25-15
25.6.3	Making InfoServer Devices Available Automatically	25-16

26 Managing the LAT Software

26.1	Understanding the LAT Protocol	26-1
26.1.1	How the LAT Protocol Works	26-2
26.1.2	Advantages of the LAT Protocol	26-2
26.2	Understanding the LAT Network	26-3
26.2.1	Service Nodes	26-4
26.2.1.1	Types of Services	26-4
26.2.1.2	Service Announcements	26-4
26.2.1.3	Print Requests	26-5
26.2.2	Terminal Server Nodes	26-5
26.2.2.1	Locating Service Nodes	26-5
26.2.2.2	Setting Up Connections	26-5
26.2.2.3	Servicing Nodes	26-5
26.2.3	Nodes Allowing Outgoing Connections	26-6
26.2.4	Components of a LAT Network	26-6
26.3	Understanding LAT Configurations	26-7
26.3.1	LAT Relationship to OpenVMS Clusters and DECnet	26-7
26.3.1.1	LAT and DECnet Running on the Same Controller	26-8
26.3.1.2	LAT and DECnet Running on Different Controllers	26-8
26.3.2	Using Multiple LAN Adapters	26-8
26.3.2.1	Supported Configurations	26-9
26.3.2.2	Unsupported Configuration	26-11
26.3.2.3	Creating Logical LAT Links	26-11
26.3.2.4	Path Discovery	26-12
26.3.2.5	Modifying LAT Parameters	26-12
26.3.3	Large Buffers in Ethernet/FDDI Configurations	26-13
26.4	Understanding the LATCP Utility	26-14
26.4.1	Invoking and Exiting LATCP	26-15
26.4.2	LATCP Commands	26-15
26.5	Starting Up the LAT Protocol	26-16
26.6	Customizing LAT Characteristics	26-18
26.6.1	Creating Additional Services	26-19
26.6.2	Setting Up Ports	26-20
26.6.2.1	Setting Up Printers	26-20
26.6.2.2	Setting Up Special Application Services	26-20
26.6.2.3	Setting Up Limited Services	26-21
26.6.3	Queuing Incoming Requests	26-21
26.6.4	Enabling Outgoing LAT Connections	26-22
26.6.5	Sample Edited LAT\$SYSTARTUP.COM Procedure	26-23
26.7	Managing the LATACP Database Size	26-24

27 Managing DECdtm Services

27.1	Understanding Transaction Logs	27-3
27.2	Planning Transaction Logs	27-3
27.2.1	Deciding the Size of a Transaction Log	27-3
27.2.2	Deciding the Location of a Transaction Log	27-4
27.3	Planning for a DECnet-Plus Network	27-4
27.3.1	Planning Your DECnet-Plus Namespace	27-4
27.3.2	Planning SCSNODE Names in Your DECnet-Plus Network	27-4
27.3.2.1	Rules for SCSNODE Names	27-5
27.3.2.2	Understanding Transaction Groups	27-5
27.4	Creating Transaction Logs	27-6
27.5	Monitoring Transaction Performance	27-8
27.6	Checking Whether a Transaction Log Is Too Small	27-11
27.7	Changing the Size of a Transaction Log	27-12
27.8	Moving a Transaction Log	27-14
27.9	Dismounting a Disk	27-17
27.10	Adding a Node	27-19
27.11	Removing a Node	27-20
27.12	Disabling DECdtm Services	27-23
27.13	Enabling DECdtm Services	27-23
27.14	Using the XA Gateway (Alpha Only)	27-24
27.14.1	Gateway Configuration	27-25

28 Managing Special Processing Environments

28.1	Understanding Multiprocessing	28-2
28.1.1	Primary and Secondary Processors	28-2
28.1.2	Available and Active Sets	28-2
28.1.3	Processor Capabilities	28-2
28.2	Managing SMP Environments	28-3
28.2.1	Creating a Multiprocessing Environment	28-3
28.2.2	Monitoring a Multiprocessing Environment	28-3
28.3	Understanding Vector Processing	28-4
28.3.1	VAX Support for Vector Processing (VAX Only)	28-4
28.3.2	VAX Vector Instruction Emulation Facility (VAX Only)	28-5
28.4	Managing the Vector Processing Environment (VAX Only)	28-5
28.4.1	Loading the Vector Processing Support Code (VAX Only)	28-5
28.4.2	Configuring a Vector Processing System (VAX Only)	28-6
28.4.3	Managing Vector Processes (VAX Only)	28-6
28.4.3.1	Adjusting System Resources and Process Quotas (VAX Only)	28-7
28.4.3.2	Distributing Scalar and Vector Resources Among Processes (VAX Only)	28-7
28.4.4	Restricting Access to the Vector Processor by Using ACLs (VAX Only)	28-8
28.4.5	Obtaining Information About a Vector Processing System (VAX Only)	28-8
28.4.5.1	DCL Lexical Functions F\$GETJPI and F\$GETSYI (VAX Only)	28-9
28.4.5.2	SHOW CPU/FULL Command (VAX Only)	28-9
28.4.5.3	SHOW PROCESS and LOGOUT/FULL Commands (VAX Only)	28-9
28.4.6	Loading the VAX Vector Instruction Emulation Facility (VVIEF) (VAX Only)	28-10

A Files–11 Disk Structure

A.1	Disk Concepts	A-1
A.1.1	Logical Organization of a Disk	A-1
A.1.2	Physical Organization of a Disk	A-2
A.2	Files–11 Structure	A-3
A.2.1	File Identification (FID)	A-4
A.2.2	ODS Directory Hierarchies	A-4
A.3	Reserved Files	A-4
A.3.1	Index File, INDEXF.SYS	A-5
A.3.1.1	Boot Block	A-6
A.3.1.2	Home Block	A-6
A.3.1.3	File Headers	A-7
A.3.2	Storage Bitmap File, BITMAP.SYS	A-8
A.3.3	Bad Block File, BADBLK.SYS	A-8
A.3.4	Master File Directory	A-8
A.3.5	Core Image File, CORIMG.SYS	A-9
A.3.6	Volume Set List File, VOLSET.SYS	A-9
A.3.7	Continuation File, CONTIN.SYS	A-9
A.3.8	Backup Log File, BACKUP.SYS	A-9
A.3.9	Pending Bad Block Log File, BADLOG.SYS	A-9
A.3.10	Quota File, QUOTA.SYS	A-9
A.3.11	Volume Security Profile, SECURITY.SYS	A-9
A.4	Files–11 ODS Level 1 (VAX Only) Versus Levels 2 and 5	A-10

B Tables of Time Differential Factors

C Compaq MIB Subagents Implemented on OpenVMS Alpha

C.1	Compaq Server MIB Subagents	C-1
C.1.1	Overview of DSM Subagents	C-2
C.1.2	Setting Up the System to Use the DSM Subagents	C-10
C.2	Compaq Cluster MIB Subagents	C-11
C.2.1	Overview of DCM Subagents	C-11
C.2.2	Setting Up the System to Use the DCM Subagents	C-13

D Managing the OpenVMS Registry

D.1	Accessing the OpenVMS Registry	D-1
D.2	Configuring the OpenVMS Registry: the REG\$CONFIG Configuration Utility	D-1
D.2.1	Configuring OpenVMS Registry Values	D-3
D.3	Starting the OpenVMS Registry	D-5
D.3.1	Starting the OpenVMS Registry Manually	D-6
D.4	Shutting Down the OpenVMS Registry	D-6
D.5	OpenVMS Registry Server Commands	D-6
	SHOW SERVER REGISTRY_SERVER	D-7
	SET SERVER REGISTRY_SERVER	D-8
D.6	OpenVMS Registry Failover in a Cluster	D-9
D.6.1	Changing the Priority of OpenVMS Registry Server Processes	D-9
D.7	Connecting to the OpenVMS Registry from a Windows NT System	D-9
D.8	OpenVMS Registry Quotas	D-10
D.9	OpenVMS Registry Security	D-10
D.10	Backing Up and Restoring the OpenVMS Registry Database	D-11

D.11	Using the OpenVMS Registry in an OpenVMS Alpha Mixed-Version Cluster	D-11
D.12	Internationalization and Unicode Support	D-11

Glossary

Index

Examples

15-1	Sample AUTOGEN Feedback Report	15-12
15-2	Sample AUTOGEN Command Procedure	15-23
20-1	Full Report Format	20-13
20-2	Brief Report Format	20-14
20-3	Terse Report Format	20-15
20-4	Summary Report Format	20-16
20-5	Fast Error (FSTERR) Report Format	20-16
20-6	Sample Operator Log File (SYS\$MANAGER:OPERATOR.LOG)	20-22
20-7	MONITOR.COM Procedure	20-42
20-8	SUBMON.COM Procedure	20-43
20-9	MONSUM.COM Procedure	20-44
D-1	Setting Priority Values	D-9
D-2	Changing Priority Values	D-9

Figures

15-1	Old and New Parameter Values	15-7
15-2	SYSMAN Temporary, Active, and Current Parameter Values	15-27
15-3	SYSGEN Temporary, Active, and Current Parameter Values	15-31
22-1	SHOW CLUSTER Default Display	22-5
22-2	SHOW CLUSTER Display with CLUSTER Report	22-9
25-1	InfoServer System Serving Clients	25-3
26-1	Components of a LAT Network	26-7
26-2	Multiple Address LAT Configuration: One LAN with Mixed-Version LAT Nodes	26-9
26-3	Multiple Address LAT Configuration: Two LANs with Mixed Version LAT Nodes	26-10
26-4	Multiple Address LAT Configuration: Two LANs with Version 5.3 LAT Nodes	26-10
26-5	Unsupported Multiple Address LAT Configuration	26-11
26-6	LAT FDDI Ring and Large Buffers	26-13
27-1	Managing DECdtm Services	27-2
27-2	Transaction Group	27-6
A-1	File Extents	A-2
A-2	Tracks and Cylinders	A-3

Tables

15-1	Controlling AUTOGEN	15-10
15-2	AUTOGEN Phases	15-16
15-3	SYSMAN PARAMETERS Commands	15-26
15-4	SYSGEN Commands Used with System Parameters	15-31
16-1	Definitions of Physical and Selective System Dumps	16-4
16-2	Paging and Swapping Terminology	16-5
16-3	Comparison of Physical and Selective System Dump Files	16-12
16-4	Symbols for Controlling the Total Size of Page, Swap, System Dump, or Error Log Dump File Space	16-31
16-5	Symbols for Controlling the Size of Individual Page and Swap Files	16-31
17-1	Libraries Recognized by LIBDECOMP.COM	17-10
17-2	Attributes of Known Images	17-18
19-1	Device Tests (VAX Only)	19-35
19-2	Device Tests (Alpha Only)	19-35
20-1	System Log Files	20-2
20-2	Parts of the Error Logging Facility	20-2
20-3	Types of Events Reported in the Error Log File	20-6
20-4	Error Log Report Options	20-8
20-5	Primary Qualifiers	20-11
20-6	(Optional) Secondary Qualifiers	20-11
20-7	Event Classes Audited by Default	20-27
20-8	Types of MONITOR Classes	20-33
20-9	MONITOR Classes	20-33
20-10	MONITOR Command Procedures	20-41
20-11	Remote Monitoring Compatibility in an OpenVMS Cluster System ...	20-45
22-1	Fields in Default SHOW CLUSTER Report	22-5
22-2	SHOW CLUSTER Qualifiers	22-8
22-3	Classes of Information Available in SHOW CLUSTER Reports	22-8
22-4	SYSMAN CONFIGURATION Commands	22-16
23-1	Choosing Compaq Networking Software for OpenVMS Systems	23-2
23-2	Compaq Networking Software Interoperability Options	23-3
23-3	Compaq TCP/IP Services for OpenVMS Documentation	23-8
23-4	DECnet-Plus for OpenVMS Documentation	23-13
24-1	Characteristics of LAN Media	24-3
24-2	LAN System Management Enhancements	24-5
24-3	LANACP System Logical Names	24-7
24-4	Functions of the LANCP Utility	24-7
24-5	Invoking the LANCP Utility	24-8
24-6	LANCP Commands	24-8
24-7	LAN Devices	24-10
24-8	SHOW DEVICE Command Qualifiers	24-11
24-9	DEFINE DEVICE and SET DEVICE Command Qualifiers	24-14
24-10	LIST DEVICE and SHOW DEVICE Command Qualifiers	24-21
24-11	DEFINE NODE and SET NODE Command Qualifiers	24-24

24-12	LIST NODE and SHOW NODE Command Qualifiers	24-29
24-13	CONNECT NODE Command Qualifiers	24-31
24-14	TRIGGER NODE Command Qualifiers	24-32
25-1	Summary of InfoServer Commands	25-7
25-2	Summary of LASTCP Commands	25-10
25-3	Summary of LADCP Commands	25-15
26-1	LATCP Commands	26-15
28-1	Settings of VECTOR_PROC System Parameter (VAX Only)	28-5
A-1	Reserved Files	A-5
A-2	Contents of Files-11 Index File	A-5
A-3	Areas of Data in File Headers	A-7
A-4	Limitations on Files-11 Structure Level 1 Volumes	A-10
B-1	TDFs for Europe	B-1
B-2	TDFs for North America	B-1
B-3	TDFs for Central and South America	B-2
B-4	TDFs for Asia	B-3
B-5	TDFs for the South Pacific	B-3
B-6	TDFs for Antarctica	B-3
C-1	DSM System Subagent Objects Implemented on OpenVMS Alpha ...	C-2
C-2	DSM Management Subagent Objects Implemented on OpenVMS Alpha	C-5
C-3	DSM Management Subagent Traps Implemented on OpenVMS Alpha	C-9
C-4	DCM Subagent Objects Implemented on OpenVMS	C-12

Preface

Intended Audience

The intended audience for this manual is Compaq *OpenVMS* system managers.

Document Structure

The *OpenVMS System Manager's Manual, Volume 2: Tuning, Monitoring, and Complex Systems* consists of the following chapters:

- Chapter 15, Managing System Parameters
- Chapter 16, Managing Page, Swap, and Dump Files
- Chapter 17, Performance Considerations
- Chapter 18, Managing File System Data Caches
- Chapter 19, Testing the System with UETP
- Chapter 20, Getting Information About the System
- Chapter 21, Tracking Resource Use
- Chapter 22, OpenVMS Cluster Considerations
- Chapter 23, Network Considerations
- Chapter 24, Managing the Local Area Network (LAN) Software
- Chapter 25, Managing InfoServer Systems
- Chapter 26, Managing the LAT Software
- Chapter 27, Managing DECdtm Services
- Chapter 28, Managing Special Processing Environments
- Appendix A, Files-11 Disk Structure
- Appendix B, Tables of Time Differential Factors
- Appendix C, Compaq MIB Subagents Implemented on OpenVMS Alpha
- Appendix D, Managing the OpenVMS Registry
- Glossary

For more information about the structure of the *OpenVMS System Manager's Manual*, see Section 1.1 and the documentation map inside the front cover of this manual.

Related Documents

The following books are helpful when used in conjunction with the *OpenVMS System Manager's Manual*:

- *OpenVMS System Management Utilities Reference Manual*
- *OpenVMS User's Manual*
- *OpenVMS Software Overview*¹
- The current version of the *Upgrade and Installation Manual* for your system
- *OpenVMS Guide to System Security*
- *OpenVMS Performance Management*
- *OpenVMS Cluster Systems and Guidelines for OpenVMS Cluster Configurations*
- The following networking manuals:
 - *Compaq TCP/IP Services for OpenVMS Installation and Configuration*
 - *Compaq TCP/IP Services for OpenVMS Management*
 - *Compaq TCP/IP Services for OpenVMS Management Command Reference*
 - *Compaq TCP/IP Services for OpenVMS Tuning and Troubleshooting*
 - *DECnet-Plus for OpenVMS Introduction and User's Guide*
 - *DECnet-Plus Planning Guide*
 - *DECnet-Plus for OpenVMS Applications Installation and Advanced Configuration*
 - *DECnet-Plus Network Control Language Reference*

For additional information about Compaq *OpenVMS* products and services, access the Compaq website at the following location:

<http://www.openvms.compaq.com/>

Reader's Comments

Compaq welcomes your comments on this manual. Please send comments to either of the following addresses:

Internet	openvmsdoc@compaq.com
Mail	Compaq Computer Corporation OSSG Documentation Group, ZKO3-4/U08 110 Spit Brook Rd. Nashua, NH 03062-2698

How to Order Additional Documentation

Visit the following World Wide Web address for information about how to order additional documentation:

<http://www.openvms.compaq.com/>

¹ This manual has been archived but is available on the *OpenVMS Documentation CD-ROM*.

Conventions

The following conventions are used in this manual:

<code>Ctrl/x</code>	A sequence such as <code>Ctrl/x</code> indicates that you must hold down the key labeled <code>Ctrl</code> while you press another key or a pointing device button.
<code>PF1 x</code>	A sequence such as <code>PF1 x</code> indicates that you must first press and release the key labeled <code>PF1</code> and then press and release another key or a pointing device button.
<code>Return</code>	In examples, a key name enclosed in a box indicates that you press a key on the keyboard. (In text, a key name is not enclosed in a box.) In the HTML version of this document, this convention appears as brackets, rather than a box.
<code>...</code>	A horizontal ellipsis in examples indicates one of the following possibilities: <ul style="list-style-type: none">• Additional optional arguments in a statement have been omitted.• The preceding item or items can be repeated one or more times.• Additional parameters, values, or other information can be entered.
<code>.</code>	A vertical ellipsis indicates the omission of items from a code example or command format; the items are omitted because they are not important to the topic being discussed.
<code>()</code>	In command format descriptions, parentheses indicate that you must enclose choices in parentheses if you specify more than one.
<code>[]</code>	In command format descriptions, brackets indicate optional choices. You can choose one or more items or no items. Do not type the brackets on the command line. However, you must include the brackets in the syntax for OpenVMS directory specifications and for a substring specification in an assignment statement.
<code> </code>	In command format descriptions, vertical bars separate choices within brackets or braces. Within brackets, the choices are optional; within braces, at least one choice is required. Do not type the vertical bars on the command line.
<code>{ }</code>	In command format descriptions, braces indicate required choices; you must choose at least one of the items listed. Do not type the braces on the command line.
bold text	This typeface represents the introduction of a new term. It also represents the name of an argument, an attribute, or a reason.
<i>italic text</i>	Italic text indicates important information, complete titles of manuals, or variables. Variables include information that varies in system output (Internal error <i>number</i>), in command lines (<i>/PRODUCER=name</i>), and in command parameters in text (where <i>dd</i> represents the predefined code for the device type).
UPPERCASE TEXT	Uppercase text indicates a command, the name of a routine, the name of a file, or the abbreviation for a system privilege.

Monospace text

Monospace type indicates code examples and interactive screen displays.

In the C programming language, monospace type in text identifies the following elements: keywords, the names of independently compiled external functions and files, syntax summaries, and references to variables or identifiers introduced in an example.

-

A hyphen at the end of a command format description, command line, or code line indicates that the command or statement continues on the following line.

numbers

All numbers in text are assumed to be decimal unless otherwise noted. Nondecimal radices—binary, octal, or hexadecimal—are explicitly indicated.

Managing System Parameters

When your system is installed or upgraded, values of system parameters are automatically set by the command procedure `SYS$UPDATE:AUTOGEN.COM` (AUTOGEN), which is supplied by Compaq. Compaq recommends that you use AUTOGEN regularly to adjust the values for system parameters to fit your hardware configuration and your system's work load.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Converting your customized parameter settings for use with AUTOGEN	Section 15.3
Modifying system parameter values with AUTOGEN (recommended method)	Section 15.5
Controlling AUTOGEN's parameter settings with MODPARAMS.DAT	Section 15.5.1
Automating AUTOGEN reports	Section 15.6
Managing system parameters with SYSMAN	Section 15.7
Managing system parameters with SYSGEN	Section 15.8
Managing system parameters with a conversational boot	Section 15.9

This chapter explains the following concepts:

Concept	Section
System parameters	Section 15.1
Default, current, and active values of system parameters	Section 15.1.1
Pages and pagelets	Section 15.1.2
The recommended method for changing system parameter values	Section 15.2
The AUTOGEN command procedure	Section 15.4
AUTOGEN feedback	Section 15.4.1
The AUTOGEN feedback report (AGEN\$PARAMS.REPORT)	Section 15.4.2
AUTOGEN phases	Section 15.4.3
The AUTOGEN parameter file (MODPARAMS.DAT)	Section 15.4.4

Managing System Parameters

15.1 Understanding System Parameters

15.1 Understanding System Parameters

The system uses values for **system parameters** to control how the system functions. System parameters control a wide range of system functions, including but not limited to the following functions:

- Memory management
- Scheduling
- Security attributes
- System caches
- Windowing system choice
- Terminal configuration
- VAXcluster or OpenVMS Cluster system attributes

The *OpenVMS System Management Utilities Reference Manual* lists and describes each system parameter.

Your distribution kit provides **default values** for system parameters to allow you to boot any supported configuration. When your system is installed or upgraded, the `SYS$UPDATE:AUTOGEN.COM` command procedure executes to evaluate your hardware configuration, estimate typical work loads, and adjust the values of system parameters as needed.

Each system parameter has associated minimum and maximum values that define the scope of allowable values.

Parameter Types

System parameters can be one or more of the following types:

Type	Description
Dynamic	The value of a dynamic system parameter can be modified while the system is active by changing the <i>active</i> value in memory. In contrast, if you change the value of a parameter that is not dynamic, you must change the <i>current</i> value stored in the parameter file, and you must reboot the system for the changed value to take effect. For information about active and current values, see Section 15.1.1.
General	The value of a general parameter affects the creation and initialization of data structures at boot time.
Major	Major parameters are most likely to require modification.
Special	Special parameters are intended for use only by Compaq. Change these parameters only if recommended by Compaq personnel or in the installation guide or release notes of a Compaq-supplied layered product.

Parameter Categories by Function

System parameters can be divided into the following categories, according to their function:

Managing System Parameters

15.1 Understanding System Parameters

Category	Function
ACP	Parameters associated with file system caches and Files-11 XQP (extended QIO procedure) or ancillary control processes (ACPs). ¹
Cluster	Parameters that affect VAXcluster or OpenVMS Cluster system operation.
Job	Parameters that control jobs.
LGI	Parameters that affect login security.
Multiprocessing	Parameters associated with symmetric multiprocessing.
PQL	Parameters associated with process creation limits and quotas.
RMS	Parameters associated with OpenVMS Record Management Services (RMS).
SCS	Parameters that control system communication services (SCS) and port driver operation. The parameters that affect SCS operation have the prefix SCS.
SYS	Parameters that affect overall system operation.
TTY	Parameters associated with terminal behavior.
User-defined	The following parameters can be user-defined: <div style="margin-left: 40px;"> USERD1 (dynamic) USERD2 (dynamic) USER3 USER4 </div>

¹Many ACP parameters are applicable only when Files-11 On-Disk Structure Level 1 disks are mounted or when an ACP is specifically requested during a mount command. In versions of the operating system before VAX VMS Version 4.0, a separate process, the ancillary control process (ACP), performed file operations such as file opens, closes, and window turns. VAX VMS Version 4.0 introduced the XQP (extended QIO procedure), which allows every process on the system to perform these operations. For compatibility reasons, the names of the parameters have not changed.

15.1.1 Default, Current, and Active Values

A system has several different sets of values for system parameters. The following table describes these values:

Value	Description
Default values	Values provided with the system to allow you to boot any supported configuration.
Current values	Values stored in the default parameter file on disk and used to boot the system. On VAX systems, the default parameter file is VAXVMSSYS.PAR. On Alpha systems, the default parameter file is ALPHAVMSSYS.PAR.
Active values	Values that are stored in memory and are used while the system is running. You can change the active value on a running system only for system parameters categorized as dynamic system parameters.
Values stored in other parameter files	For special purposes, you can create a parameter file other than the default parameter file that is used to store current values.

Managing System Parameters

15.1 Understanding System Parameters

When the system boots, it reads the current values into memory, creating active values. An active value remains equal to the current value until you change either the active value or the current value.

When you execute the AUTOGEN command procedure through the SETPARAMS phase, it changes *current* values.

The System Management utility (SYSMAN) and the System Generation utility (SYSGEN) allow you to show and modify both *current* and *active* values. Use the USE and WRITE commands to specify which values you want to show or modify.

For more information about managing parameters with SYSMAN, see Section 15.7. For more information about managing parameters with SYSGEN, see Section 15.8.

15.1.2 Pages and Pagelets

On VAX systems, the operating system allocates and deallocates memory for processes in units called **pages**. A page on a VAX system is 512 bytes. Some system parameter values are allocated in units of pages.

On Alpha systems, some system parameter values are allocated in units of pages, while others are allocated in units of **pagelets**.

A page on an Alpha system can be 8 kilobytes (KB) (8192 bytes), 16KB, 32KB, or 64KB. A pagelet is a 512-byte unit of memory. One Alpha pagelet is the same size as one VAX page. On an Alpha computer with a page size of 8KB, 16 Alpha pagelets equal one Alpha page.

When reviewing parameter values, especially those parameters related to memory management, be sure to note the units required for each parameter. Section 15.7.2 and Section 15.8.2 explain how to show parameter values and their units of allocation.

15.2 Recommended Method for Changing Parameter Values

Many system parameters can affect other parameters and the performance of the system. For this reason, Compaq recommends that you use the command procedure SYS\$UPDATE:AUTOGEN.COM (AUTOGEN) to manage system parameters. For information about AUTOGEN, see Section 15.4.

The System Management utility (SYSMAN) and the System Generation utility (SYSGEN) also allow you to manage system parameters. Although these utilities are not generally recommended for *changing* parameter values, you can use one of these utilities for the following reasons:

- To display system parameters and their values on a VAX or Alpha system
- To display system parameters and their values for systems in an OpenVMS Cluster environment
- To temporarily modify a single parameter that has little effect on other parameters

Caution

If you change a parameter value with SYSMAN or SYSGEN, the value you set will be overridden or reset to the default value when you run AUTOGEN. To retain the changes when you run AUTOGEN, you must add the parameter value to the AUTOGEN parameter file MODPARAMS.DAT. For more information, see Section 15.5.1.

Managing System Parameters

15.2 Recommended Method for Changing Parameter Values

If you currently use SYSMAN or SYSGEN to change parameters, and you have not added your customized parameter settings to MODPARAMS.DAT, follow the instructions in Section 15.3 before running AUTOGEN.

15.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

Compaq recommends that you use the AUTOGEN command procedure to tune your system. If you use the System Management utility (SYSMAN) or the System Generation utility (SYSGEN) to modify system parameter values, and you do not include these changes in the AUTOGEN parameter file MODPARAMS.DAT, these changes will be overridden the next time you run AUTOGEN.

If you used SYSMAN or SYSGEN to change parameter values in the past, use the following procedure to convert your parameter settings to work with AUTOGEN. This procedure explains how to add your customized parameter settings to MODPARAMS.DAT so they will be retained when you run AUTOGEN.

Before performing this task, you should understand AUTOGEN, feedback, and the AUTOGEN parameter file MODPARAMS.DAT, as explained in Section 15.4.

1. Save the parameter values that the system is now using as follows:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE ACTIVE
SYSMAN> PARAMETERS WRITE SYS$SYSTEM:nodename_PARAMS_CURRENT.PAR
```

2. Write a listing of the active parameter values to an ASCII file named *nodename_PARAMS.OLD* as follows:

```
SYSMAN> PARAMETERS SHOW/ALL/OUTPUT=nodename_PARAMS.OLD
SYSMAN> PARAMETERS SHOW/SPECIAL/OUTPUT=nodename_PARAMS_SPECIAL.OLD
SYSMAN> EXIT
$ APPEND nodename_PARAMS_SPECIAL.OLD nodename_PARAMS.OLD
```

You will use this file in step 6.

3. Edit AUTOGEN's parameter file SYS\$SYSTEM:MODPARAMS.DAT to define symbols to specify values for the following parameters:

- Parameter values that are not calculated by AUTOGEN, such as SCSNODE and SCSSYSTEMID. Refer to the AUTOGEN description in the *OpenVMS System Management Utilities Reference Manual* for a table of the parameters calculated by AUTOGEN.
- Any parameter values that must be adjusted to suit your system work load, for example, GBLPAGES and GBLSECTIONS.

To specify a value, define symbols using the format MIN_parameter, MAX_parameter, or ADD_parameter rather than specifying an explicit value. For example:

```
$ EDIT SYS$SYSTEM:MODPARAMS.DAT

SCSNODE = "MYNODE" ! Not calculated by AUTOGEN
SCSSYSTEMID = 10001 ! Not calculated by AUTOGEN
MIN_GBLPAGES = 10000 ! Needed for MCS, BLISS32, and ADA
MIN_GBLSECTIONS = 600 ! Needed for MCS, BLISS32, and ADA
```

Managing System Parameters

15.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

To help you track the changes you make in MODPARAMS.DAT, add comments to each line, preceded by an exclamation point (!). For information about defining symbols in MODPARAMS.DAT, see Section 15.5.1.

4. Run AUTOGEN, but do *not* reboot. Use one of the following commands, depending on your system:

- If the system has run a typical work load for more than 24 hours since last booting:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS SETPARAMS FEEDBACK
```

The SAVPARAMS phase collects feedback information about resource use on the running system; this information is used by AUTOGEN. This command creates a feedback report named SYS\$SYSTEM:AGEN\$PARAMS.REPORT, which tells you about peak resource use.

- If you want to use a previously collected feedback file:

```
$ @SYS$UPDATE:AUTOGEN GETDATA SETPARAMS FEEDBACK
```

If you start from the GETDATA phase, AUTOGEN does not collect current feedback.

- If this is a new system (that is, it has no feedback) or the system has had little activity since last boot (for example, over the weekend) so there is no valid feedback file:

```
$@SYS$UPDATE:AUTOGEN GETDATA SETPARAMS CHECK_FEEDBACK
```

Use the CHECK_FEEDBACK parameter to let AUTOGEN determine whether the feedback is valid.

5. Write a listing of the new parameter values to an ASCII file as follows:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE CURRENT
SYSMAN> PARAMETERS SHOW /ALL /OUTPUT=nodename_PARAMS.NEW
SYSMAN> PARAMETERS SHOW /SPECIAL /OUTPUT=nodename_PARAMS_SPECIAL.NEW
SYSMAN> EXIT
$ APPEND nodename_PARAMS_SPECIAL.NEW nodename_PARAMS.NEW
```

6. Compare the old and new parameter values as follows:

```
$ DIFFERENCES/PARALLEL/OUTPUT=nodename_PARAMS.DIF/MATCH=5 -
_$ nodename_PARAMS.OLD nodename_PARAMS.NEW
```

7. Print the differences file you created in step 6 (named in the format *nodename_PARAMS.DIF*). Print the file on a 132-column line printer to make the output easier to read.
8. Compare the numbers in the two columns following each parameter name column. The left column shows the old value; the right column shows the new value. Figure 15–1 illustrates sample output.

Managing System Parameters

15.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

Figure 15–1 Old and New Parameter Values

Parameter Names	1					1				
	6	10000	512	-	-	6	10000	512	-	-
GBLPAGES	77500					81800				
SYSMWCNT	2400	500	40	1638		2800	500	40	1638	
INTSTKPAGES	4	4	1	-		4	4	1	-	
BALSETCNT	250	16	4	819		250	16	4	819	
WSMAX	32800	1024	60	20000		65500	1024	60	20000	
NPAGEDYN	1944576	360000	16384	-		3000000	360000	16384	-	
NPAGEVIR	777328	1000000	16384	-		12000000	1000000	16384	-	
PAGEDYN	1516032	190000	10240	-		1780056	190000	10240	-	
VIRTUALPAGECNT	150000	9216	512	100000		270144	9216	512	100000	
⋮	⋮					⋮				
	⋮	Old Values				⋮	New Values			

ZK-5175A-GE

9. Make any adjustments in MODPARAMS.DAT using symbols prefixed by MIN_, MAX_, or ADD_. For example, if AUTOGEN calculated a smaller value for GBLPAGES, you might specify a minimum value for this parameter as follows:

```
MIN_GBLPAGES = 10000
```

If you originally specified a parameter value in MODPARAMS.DAT (in step 3) but the parameter has not been changed, verify the following data:

- The parameter name is spelled correctly and completely (not abbreviated). In MODPARAMS.DAT, AUTOGEN sees parameter names as symbol assignments. AUTOGEN cannot equate a symbol to the corresponding system parameter unless it is spelled correctly. Look in AGEN\$FEEDBACK.REPORT for any error messages AUTOGEN might have written.
- The value is correct: count the digits and make sure no commas are present.
- The parameter occurs only once in MODPARAMS.DAT.
- The parameter is not commented out.

For most parameters, if the new value is greater than the old value, you can accept AUTOGEN's setting. If the new value is less than the old value, Compaq recommends that you retain the old value because the system may not have been using that resource when running AUTOGEN.

For example, you might have used SYSMAN to increase GBLPAGES to 10,000 to accommodate layered products, but have not specified that change in MODPARAMS.DAT. AUTOGEN might calculate that the system needs only 5000 global pages. When you reboot after running AUTOGEN, not all of your layered products may be installed, and you might receive the system message GPTFULL, "global page table full," indicating that the system needs more GBLPAGES.

10. Repeat from step 3 until you are satisfied with the new parameter values. If necessary, make further changes in MODPARAMS.DAT, run AUTOGEN again, and verify the changes as before. Usually after this second pass of AUTOGEN, the parameter values will be stable and you can then reboot.

Managing System Parameters

15.3 Converting Your Customized Parameter Settings for Use with AUTOGEN

11. Reboot. When you reboot, the system will use the new parameter values. Using AUTOGEN to reboot or rebooting right away is not necessary. However you must reboot before the system uses the new parameter values.

If the system does not boot, perform a conversational boot and use the backup parameter file you created in step 1:

```
SYSBOOT> USE SYS$SYSTEM:nodename_PARAMS_CURRENT.PAR
SYSBOOT> CONTINUE
```

When you enter the CONTINUE command, the system boots with the parameter values you saved before running AUTOGEN.

After the system has booted, if you want to use the old parameter values you can enter the following commands:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE SYS$SYSTEM:nodename_PARAMS_CURRENT.PAR
SYSMAN> PARAMETERS WRITE CURRENT
SYSMAN> EXIT
```

12. Run AUTOGEN using feedback regularly to ensure that the resources of your system match your system work load. For information about running AUTOGEN using feedback, see Section 15.5.

15.4 Understanding the AUTOGEN Command Procedure

The AUTOGEN command procedure, SYS\$UPDATE:AUTOGEN.COM, is provided on your distribution kit, and runs automatically when your system is installed or upgraded to set appropriate values for system parameters. In addition, Compaq recommends that you run AUTOGEN when you want to reset values for system parameters or to resize page, swap, and dump files. The new values and file sizes take effect the next time the system boots.

AUTOGEN only calculates certain significant system parameters. Refer to the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual* for a table of system parameters affected by AUTOGEN calculation.

When to Run AUTOGEN

Compaq recommends running AUTOGEN in the following circumstances:

- During a new installation or upgrade. (This happens automatically as part of the installation or upgrade procedure.)
- Whenever your work load changes significantly.
- When you add an optional (layered) software product. See the specific product documentation for installation requirements. Certain layered products might require you to execute AUTOGEN to adjust parameter values and page and swap file sizes. (For information about using AUTOGEN to modify page and swap files, see Section 16.16.1.)
- When you install images with the /SHARED attribute; the GBLSECTIONS and GBLPAGES parameters might need to be increased to accommodate the additional global pages and global sections consumed.
- On a regular basis to monitor changes in your system's work load. You can automate AUTOGEN to regularly check feedback and recommend system parameter changes. Section 15.6 describes a batch-oriented command procedure that runs AUTOGEN in feedback mode on a regular basis and automatically sends the feedback report to an appropriate Mail account.

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

- Periodically to provide adequate swapping file space. Use the FEEDBACK option and make sure the system has been up long enough (at least 24 hours) and that the load is typical. Also, make sure the SYS\$SYSTEM:MODPARAMS.DAT file does not contain a hardcoded SWAPFILE value, which prevents AUTOGEN from correctly sizing the swapping files.

AUTOGEN Operations

AUTOGEN executes in phases. Depending on which phases you direct it to execute, AUTOGEN performs some or all of the following operations:

- Collects the following types of data:
 - Feedback (from the running system)
 - The hardware configuration (from the system)
 - Parameter requirements supplied by you (from MODPARAMS.DAT)
 - Parameter requirements supplied by Compaq
- Calculates appropriate new values for significant system parameters (listed in the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual*)
- Creates a new installed image list
- Calculates the sizes of system page, swap, and dump files
- Adjusts the sizes of system page, swap, and dump files values of system parameter values, if necessary
- Optionally shuts down and reboots the system

Invoking AUTOGEN

To invoke AUTOGEN, enter a command in the following format at the DCL prompt:

```
@SYS$UPDATE:AUTOGEN [start-phase] [end-phase] [execution-mode]
```

where:

<i>start-phase</i>	Is the phase where AUTOGEN is to begin executing. Section 15.4.3 lists the AUTOGEN phases.
<i>end-phase</i>	Is the phase where AUTOGEN is to complete executing. Section 15.4.3 lists the AUTOGEN phases.
<i>execution-mode</i>	Is one of the following modes: <ul style="list-style-type: none">• FEEDBACK—Use feedback.• NOFEEDBACK—Do not use feedback.• CHECK_FEEDBACK—Use feedback if it is valid. If feedback is invalid, ignore it, but continue executing through the end phase.• Blank (no execution mode specified)—Use feedback if it is valid. If it is not valid, quit before making any modifications.

For detailed information about invoking AUTOGEN, and the command line parameters you can specify, refer to the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual*.

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

Controlling AUTOGEN Operations

Table 15–1 summarizes the methods for controlling AUTOGEN behavior.

Table 15–1 Controlling AUTOGEN

To Control...	Use This Method...
Which operations AUTOGEN is to perform	Specify a start phase and an end phase when you invoke AUTOGEN.
Parameter values set by AUTOGEN	Specify values in the AUTOGEN parameter file MODPARAMS.DAT. Periodically examine the results of calculations that AUTOGEN makes to determine whether AUTOGEN has drawn the correct conclusions about your hardware configuration and to be sure the system parameter values are appropriate for your workload requirements. If the values are not appropriate, adjust them by specifying desired values in MODPARAMS.DAT. For more information on MODPARAMS.DAT, see Section 15.4.4.
AUTOGEN's use of feedback information	Specify an execution mode when you invoke AUTOGEN. AUTOGEN can often improve system performance by using dynamic feedback gathered from the running system. However, feedback information is not always valid or appropriate. For more information, see Section 15.4.1.

15.4.1 AUTOGEN Feedback

AUTOGEN feedback minimizes the need for you to modify parameter values or system file sizes. Instead, feedback allows AUTOGEN to automatically size the operating system based on your actual work load. **Sizing** is the process of matching the allocation of system resources (memory and disk space) with the workload requirements of your site.

Feedback is information, continuously collected by the operating system executive, about the amount of various resources the system uses to process its work load. The information is collected when exception events occur, so the collection does not affect system performance. When run in *feedback mode*, AUTOGEN analyzes this information and adjusts any related parameter values.

Note

When running AUTOGEN after making a major configuration change, specify **nofeedback** to assure the use of initial AUTOGEN settings. See Section 15.4.

AUTOGEN feedback affects the following resources (for a complete list of the affected system parameters, refer to the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual*):

- Nonpaged pool
- Paged pool
- Lock resources
- Number of processes
- Global pages

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

- Global sections
- File system caches
- System logical name table sizes
- Page files
- Swap files

Feedback is gathered during AUTOGEN's SAVPARAMS phase and is written to the file `SYS$SYSTEM:AGEN$FEEDBACK.DAT`. This file is then read during the GETDATA phase. (See Section 15.4.3 for more information about AUTOGEN phases.)

Feedback is useful only if it accurately reflects the system's normal work load. For this reason, AUTOGEN performs some basic checks on the feedback and issues a warning message for either of the following conditions:

- The system has been up for less than 24 hours.
- The feedback is over 30 days old.

Whenever you modify the system (for example, a hardware upgrade, a change in the number of users, an optional product installation), you should operate in the new system environment for a period of time, and then execute AUTOGEN again starting from the SAVPARAMS phase.

On VAX systems, you can define the logical name `AGEN$FEEDBACK_REQ_TIME` to specify, in hours, a minimum age required for feedback. For more information, see Section 15.5.2.

When AUTOGEN runs, it displays whether feedback is used, as follows:

```
Feedback information was collected on 21-JAN-2000 14:00:08.53
Old values below are the parameter values at the time of collection.
The feedback data is based on 21 hours of up time.
Feedback information will be used in the subsequent calculations
```

15.4.2 Feedback Report (AGEN\$PARAMS.REPORT)

Decides whether to use the system parameter values and system file sizes calculated by AUTOGEN. To help in your decision making, AUTOGEN generates a report file (`SYS$SYSTEM:AGEN$PARAMS.REPORT`) that includes the following information:

- All parameters and system files directly affected by the feedback
- Current values
- New values
- The feedback used in each parameter calculation
- Any user- or Compaq-supplied modifications found in `MODPARAMS.DAT`
- Any advisory or warning messages displayed during AUTOGEN's operations
- On VAX systems, any user- or Compaq-supplied modifications found in `VMSPARAMS.DAT`
- On Alpha systems, the parameters found during the GENPARAMS phase

Example 15-1 shows the contents of a sample AUTOGEN feedback report for a VAX system. On Alpha systems, the feedback report is similar but not identical to this example.

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

Suppressing Informational Messages

To suppress the display of informational messages, define the `AGEN$REPORT_NO_INFORMATIONALS` logical to `TRUE`. Messages are entered in `SYS$SYSTEM:AGEN$PARAMS.REPORT` regardless of the value of `AGEN$REPORT_NO_INFORMATIONALS`.

Filtering DCL Statements from Your Report

The feedback report will contain DCL statements in `MODPARAMS.DAT` that are not simple assignments to system parameters or the `ADD_`, `MAX_`, or `MIN_` extensions. To filter these statements out of the report, begin each statement in `MODPARAMS.DAT` with a dollar sign (`$`).

Example 15-1 Sample AUTOGEN Feedback Report

```
AUTOGEN Parameter Calculation Report on node: NODE22
  This information was generated at 23-APR-2000 01:45:47.87
  AUTOGEN was run from GETDATA to TESTFILES using FEEDBACK

** No changes will be done by AUTOGEN **
  The values given in this report are what AUTOGEN would
  have set the parameters to.

Processing Parameter Data files
-----

** WARNING ** - The system was up for less than 24 hours when the feedback
information was recorded. This could result in feedback information
that does not accurately reflect your typical work load.

Including parameters from: SYS$SYSTEM:MODPARAMS.DAT
The following was detected within MODPARAMS.DAT
  Please review immediately.

** INFORMATIONAL ** - Multiple MIN values found for MIN_CHANNELCNT.
  Using MODPARAMS value (550) which is superseding OpenVMS value (255)

** INFORMATIONAL ** - Multiple MIN values found for MIN_SWPOUTPGCNT.
  Using MODPARAMS value (1000) which is superseding OpenVMS value (500)

** INFORMATIONAL ** - Multiple MIN values found for MIN_PQL_DWSEXTENT.
  Using MODPARAMS value (11000) which is superseding OpenVMS value (1024)

** INFORMATIONAL ** - Multiple MIN values found for MIN_PQL_MWSEXTENT.
  Using MODPARAMS value (11000) which is superseding OpenVMS value (1024)
Feedback information was collected on 22-APR-2000 14:00:07.70
  Old values below are the parameter values at the time of collection.
  The feedback data is based on 13 hours of up time.
  Feedback information will be used in the subsequent calculations

Parameter information follows:
-----

MAXPROCESSCNT parameter information:
  Feedback information.
    Old value was 100, New value is 80
    Maximum Observed Processes: 52

Information about VMS executable image Processing:
  Processing SYS$MANAGER:VMS$IMAGES_MASTER.DAT
```

(continued on next page)

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

Example 15-1 (Cont.) Sample AUTOGEN Feedback Report

GBLPAGFIL parameter information:
Override Information - parameter calculation has been overridden.
The calculated value was 1024. The new value is 6024.
GBLPAGFIL has been increased by 5000.
GBLPAGFIL is not allowed to be less than 6024.

GBLPAGES parameter information:
Feedback information.
Old value was 43300, New value is 50000
Peak used GBLPAGES: 36622
Global buffer requirements: 6024

GBLSECTIONS parameter information:
Feedback information.
Old value was 400, New value is 400
Peak used GBLSECTIONS: 294
Override Information - parameter calculation has been overridden.
The calculated value was 350. The new value is 400.
GBLSECTIONS is not allowed to be less than 400.

LOCKIDTBL parameter information:
Feedback information.
Old value was 2943, New value is 3071
Current number of locks: 1853
Peak number of locks: 3200

LOCKIDTBL_MAX parameter information:
Feedback information.
Old value was 65535, New value is 65535

RESHASHTBL parameter information:
Feedback information.
Old value was 1024, New value is 1024
Current number of resources: 957

MSCP_LOAD parameter information:
Override Information - parameter calculation has been overridden.
The calculated value was 1. The new value is 0.
MSCP_LOAD has been disabled by a hard-coded value of 0.

MSCP_BUFFER parameter information:
Feedback information.
Old value was 128, New value is 128
MSCP server I/O rate: 0 I/Os per 10 sec.
I/Os that waited for buffer space: 0
I/Os that fragmented into multiple transfers: 0

SCSCONNCNT parameter information:
Feedback information.
Old value was 5, New value is 5
Peak number of nodes: 1
Number of CDT allocation failures: 0

SCSRESPCNT parameter information:
Feedback information.
Old value was 300, New value is 300
RDT stall count: 0

SCSBUFFCNT parameter information:
Feedback information.
Old value was 512, New value is 512
CIBDT stall count: 0

(continued on next page)

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

Example 15–1 (Cont.) Sample AUTOGEN Feedback Report

NPAGEDYN parameter information:
Feedback information.
Old value was 686592, New value is 783360
Maximum observed non-paged pool size: 815616 bytes.
Non-paged pool request rate: 47 requests per 10 sec.

LNMSHASHTBL parameter information:
Feedback information.
Old value was 1024, New value is 1024
Current number of shareable logical names: 1194

ACP_DIRCACHE parameter information:
Feedback information.
Old value was 88, New value is 88
Hit percentage: 99%
Attempt rate: 0 attempts per 10 sec.

ACP_DINDXCACHE parameter information:
Feedback information.
Old value was 25, New value is 25
Hit percentage: 97%
Attempt rate: 1 attempts per 10 sec.

ACP_HDRCACHE parameter information:
Feedback information.
Old value was 88, New value is 106
Hit percentage: 98%
Attempt rate: 17 attempts per 10 sec.

ACP_MAPCACHE parameter information:
Feedback information.
Old value was 8, New value is 8
Hit percentage: 2%
Attempt rate: 4 attempts per 10 sec.

PAGEDYN parameter information:
Feedback information.
Old value was 521728, New value is 542208
Current paged pool usage: 304160 bytes.
Paged pool request rate: 1 requests per 10 sec.

PFRATL parameter information:
Override Information - parameter calculation has been overridden.
The calculated value was 0. The new value is 1.
PFRATL has been disabled by a hard-coded value of 1.

WSDEC parameter information:
Override Information - parameter calculation has been overridden.
The calculated value was 35. The new value is 19.
WSDEC has been disabled by a hard-coded value of 19.

MPW_LOLIMIT parameter information:
Override Information - parameter calculation has been overridden.
The calculated value was 120. The new value is 2100.
MPW_LOLIMIT is not allowed to be less than 2100.

MPW_HILIMIT parameter information:
Override Information - parameter calculation has been overridden.
The calculated value was 1310. The new value is 4500.
MPW_HILIMIT is not allowed to be less than 4500.

(continued on next page)

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

Example 15-1 (Cont.) Sample AUTOGEN Feedback Report

LONGWAIT parameter information:

Override Information - parameter calculation has been overridden.
The calculated value was 30. The new value is 10.
LONGWAIT has been disabled by a hard-coded value of 10.

WSMAX parameter information:

Override Information - parameter calculation has been overridden.
The calculated value was 8200. The new value is 12000.
WSMAX is not allowed to be less than 12000.

PROCSECTCNT parameter information:

Override Information - parameter calculation has been overridden.
The calculated value was 32. The new value is 40.
PROCSECTCNT is not allowed to be less than 40.

PQL_DWSEXTENT parameter information:

Override Information - parameter calculation has been overridden.
The calculated value was 400. The new value is 11000.
PQL_DWSEXTENT is not allowed to be less than 11000.

PQL_MWSEXTENT parameter information:

Override Information - parameter calculation has been overridden.
The calculated value was 2048. The new value is 11000.
PQL_MWSEXTENT is not allowed to be less than 11000.

VAXCLUSTER parameter information:

Override Information - parameter calculation has been overridden.
The calculated value was 1. The new value is 0.
VAXCLUSTER has been disabled by a hard-coded value of 0.

Page, Swap, and Dump file calculations

Page and Swap file calculations.

PAGEFILE1_SIZE parameter information:

Feedback information.
Old value was 45200, New value is 50500
Maximum observed usage: 25265
PAGEFILE1_SIZE will be modified to hold 50500 blocks

PAGEFILE2_SIZE parameter information:

Feedback information.
Old value was 154000, New value is 194400
Maximum observed usage: 97175
PAGEFILE2_SIZE will be modified to hold 194400 blocks

** WARNING ** - The disk on which PAGEFILE2 resides would be
over 95% full if it were modified to hold 194400 blocks.
NODE22\$DKA300:[SYSTEM_FILES]PAGEFILE.SYS will not be modified.
NODE22\$DKA300:[SYSTEM_FILES]PAGEFILE.SYS will remain at 154002 blocks.

SWAPFILE1_SIZE parameter information:

Feedback information.
Old value was 15000, New value is 15000
Maximum observed usage: 14280
Override Information - parameter calculation has been overridden.
The calculated value was 21400. The new value is 15000.
SWAPFILE1 SIZE is not allowed to exceed 15000.
SWAPFILE1 will not be modified.

(continued on next page)

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

Example 15–1 (Cont.) Sample AUTOGEN Feedback Report

```
SWAPFILE2_SIZE parameter information:
Feedback information.
  Old value was 50000, New value is 26300
  Maximum observed usage: 1680
SWAPFILE2_SIZE will be modified to hold 26300 blocks

** WARNING ** - The disk on which SWAPFILE2 resides would be
over 95% full if it were modified to hold 26300 blocks.
NODE22$DKA300:[SYSTEM_FILES]SWAPFILE.SYS will not be modified.
NODE22$DKA300:[SYSTEM_FILES]SWAPFILE.SYS will remain at 50001 blocks.

Dumpfile calculations:

No dump file modifications would have been made.
Dumpfile will remain at 34116 blocks.
```

15.4.3 AUTOGEN Phases

When you invoke AUTOGEN, you specify a start phase and an end phase for AUTOGEN to execute. AUTOGEN executes all phases from the start phase to the end phase. Depending on the start phase and end phase you specify, AUTOGEN can execute any of the following phases, in the order shown in Table 15–2.

Table 15–2 AUTOGEN Phases

Phase	Description
SAVPARAMS	Saves dynamic feedback from the running system.
GETDATA	Collects all data to be used in AUTOGEN calculations.
GENPARAMS	Generates new system parameters; creates the installed image list.
TESTFILES	Displays the system page, swap, and dump file sizes calculated by AUTOGEN (cannot be used as a start phase).
GENFILES	Generates new system page, swap, and dump files if appropriate (cannot be used as a start phase).
SETPARAMS	Runs SYSMAN to set the new system parameters in the default parameter file, saves the original parameters, and generates a new parameter file, AUTOGEN.PAR. On VAX systems, the default parameter file is VAXVMSSYS.PAR. The original parameters are saved in the file VAXVMSSYS.OLD. On Alpha systems, the default parameter file is ALPHAVMSSYS.PAR. The original parameters are saved in the file ALPHAVMSSYS.OLD.
SHUTDOWN	Prepares the system to await a manual reboot.
REBOOT	Automatically shuts down and reboots the system.
HELP	Displays help information to the screen.

For detailed information about each AUTOGEN phase and the files affected by each phase, refer to the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual*.

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

15.4.4 AUTOGEN Parameter File (MODPARAMS.DAT)

AUTOGEN reads a parameter file named MODPARAMS.DAT during the GETDATA phase. You can add commands to this file to control the system parameter values and file sizes that AUTOGEN sets. You can use MODPARAMS.DAT to perform the following actions:

Operation	For More Information
Increase the value of any numeric system parameter	Section 15.5.1.1
Set a minimum value for a numeric system parameter	Section 15.5.1.2
Set a maximum value for a numeric system parameter	Section 15.5.1.3
Specify an absolute value for a system parameter	Section 15.5.1.4
Include an external parameter file	Section 15.5.3
Specify sizes of page, swap, and dump files	Section 16.16.1.2
†Define the number of VAXcluster nodes	Section 15.5.1.5
†Define the number of Ethernet adapters	Section 15.5.1.6
†Preset parameter values before adding memory	Section 15.5.1.7
Specify an alternate default startup command procedure	Section 4.5.2

†VAX specific

To help track changes you make to MODPARAMS.DAT, make sure you add comments, preceded by the exclamation point (!), each time you change the file.

Caution

The recommended method of changing system parameters and system file sizes is to edit MODPARAMS.DAT to include parameter settings. If you change a system parameter value or file size using SYSMAN, SYSGEN, or a conversational boot, and you do not specify the value in MODPARAMS.DAT, AUTOGEN will recalculate the value or file size the next time it runs. For more information, see Section 15.5.1.

Example

The following example shows the contents of a sample MODPARAMS.DAT file:

```
!
! ***** A Sample MODPARAMS.DAT for Node NODE22 *****
!
! MODPARAMS.DAT for "NODE22"
! REVISED: 04/29/00 -CHG- Upped GBLPAGES to account for ADA.
!
SCSNODE          = "NODE22"      ! This is not calculated by AUTOGEN.
SCSSYSTEMID      = 19577         ! This is not calculated by AUTOGEN.
TTY_DEFCHAR2     = %X0D34       ! This is not calculated by AUTOGEN.
ADD_ACP_DIRCACHE= 150           ! Hit rate was only 65% on directory cache.
MIN_PAGEDYN      = 500000       ! PAGEDYN must be at least 1/2 Mbyte to
                                ! account for a large number of logical names.
!
MAX_PAGEFILE1_SIZE = 15000      ! Maximum size for primary page.
MAX_SWAPFILE      = 5000        ! Maximum size for swap file space.
MAX_DUMPFIL      = 32768       ! Maximum size for dump file space.
```

Managing System Parameters

15.4 Understanding the AUTOGEN Command Procedure

```
ADD_GBLPAGES      = 425+507+157    ! Account for MCS, BLISS32 and ADA.
ADD_GBLSECTIONS  = 4 + 5 + 2      ! Account for MCS, BLISS32 and ADA.
VIRTUALPAGECNT   = 144264         ! So that we can read MONSTR's 68Mb dumps.
!
! end of MODPARAMS.DAT for NODE22
```

15.5 Modifying System Parameters with AUTOGEN

The recommended method of modifying system parameters is to execute AUTOGEN in two passes, as follows:

1. First pass—Execute AUTOGEN using the following command:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS TESTFILES
```

This command instructs AUTOGEN to perform the following actions:

- Save the current feedback
- Gather all the information required for the calculations
- Calculate the system parameter values
- Generate the feedback report
- Write the information to SETPARAMS.DAT
- Compute the sizes of page, swap, and dump files; log the sizes in the report file AGEN\$PARAMS.REPORT.

Review the input to the calculations (PARAMS.DAT), the output from the calculations (SETPARAMS.DAT), and the report generated (AGEN\$PARAMS.REPORT).

If you are not satisfied with the parameter settings, modify the parameter values by editing MODPARAMS.DAT as explained in Section 15.5.1. If you are not satisfied with the file sizes, modify the sizes as explained in Section 16.16. Then reexecute AUTOGEN from the GETDATA phase.

When you are satisfied with the contents of SETPARAMS.DAT, go to step 2.

2. Second pass—Execute AUTOGEN a second time using the following command:

```
$ @SYS$UPDATE:AUTOGEN GENPARAMS REBOOT
```

This AUTOGEN command runs SYSMAN to update the new system parameter values and installs them on the system when it is rebooted.

15.5.1 Controlling AUTOGEN's Parameter Settings with MODPARAMS.DAT

If, after examining the AGEN\$PARAMS.REPORT or SETPARAMS.DAT file, you decide to correct hardware configuration data or modify system parameter values chosen by AUTOGEN, edit the MODPARAMS.DAT file as described in this section to manually specify parameter values.

Caution

Always edit MODPARAMS.DAT to specify values for parameters. Do not edit PARAMS.DAT; modifying the contents of this file might prevent AUTOGEN from operating correctly.

Managing System Parameters

15.5 Modifying System Parameters with AUTOGEN

For information about editing MODPARAMS.DAT to control sizes of page, swap, and dump files, see Section 16.16.1.2.

You can define symbols in MODPARAMS.DAT using the following formats to control parameter values:

Control Method	Symbol Format	For More Information
Increase a value by a specified amount	ADD_*	Section 15.5.1.1
Specify a minimum value	MIN_*	Section 15.5.1.2
Specify a maximum value	MAX_*	Section 15.5.1.3
Specify an absolute value	Parameter name	Section 15.5.1.4

When defining symbols in MODPARAMS.DAT, make sure to adhere to the following criteria:

- The value is correct and valid for the parameter. Count the digits. Do not use commas.
- The symbol occurs only once in MODPARAMS.DAT.
- The symbol value is not commented out.
- The symbol name is spelled correctly and completely (not abbreviated).
- String values must be within quotation marks ("").

Caution

When AUTOGEN reads MODPARAMS.DAT or any other parameter file, it checks to determine if the symbol names specified in the file are valid. If they are not, AUTOGEN writes a warning message to AGEN\$PARAMS.REPORT. However, AUTOGEN checks only the symbol name; it does not check the validity of the value specified for the symbol.

If a value is invalid, the line is *not* ignored. AUTOGEN attempts to use the specified value.

A symbol is not checked if it is specified in a line that contains a DCL expression other than the equal sign (=). For example, AUTOGEN does not check the validity of a symbol name specified in a line with the DCL IF statement. Instead, AUTOGEN writes a warning message to AGEN\$PARAMS.REPORT.

To help track changes you make to MODPARAMS.DAT, make sure you add comments preceded by an exclamation point (!) each time you change the file.

15.5.1.1 Increasing a Value with the ADD_ Prefix

Use the ADD_ prefix to increase the value of any numeric parameter. The new values are updated in subsequent AUTOGEN calculations during the GENPARAMS phase. The following example demonstrates the use of the ADD_ prefix:

```
ADD_GBLPAGES=500
ADD_NPAGEDYN=10000
```

An ADD_ parameter record for a parameter that AUTOGEN calculates will add the value to AUTOGEN's calculations. An ADD_ parameter record for

Managing System Parameters

15.5 Modifying System Parameters with AUTOGEN

a parameter that AUTOGEN does not calculate will add the value to the parameter's default (not current) value. (Refer to the AUTOGEN section on the *OpenVMS System Management Utilities Reference Manual* for a table of parameters affected by AUTOGEN.)

Note

The ADD_ value is added to the calculated value once, and does not accumulate with successive runs for feedback calculations.

Typically, you would not use the ADD_ prefix for modifying parameters that are calculated by the feedback mechanism, because the feedback results should accurately reflect your work load. However, if you do use the ADD_ prefix with feedback, AUTOGEN adds a value only once if AUTOGEN is run to the SETPARAMS phase or beyond. To maintain a minimum level above AUTOGEN's calculation, use the MIN_ prefix.

15.5.1.2 Specifying a Minimum Value with the MIN_ Prefix

Use the MIN_ prefix if you do not want AUTOGEN to set a parameter below a specified value. MIN_ refers to the minimum value to which a parameter can be set by AUTOGEN. The following example sets the minimum value to 400,000:

```
MIN_PAGEDYN = 400000
```

15.5.1.3 Specifying a Maximum Value with the MAX_ Prefix

Use the MAX_ prefix if you do not want AUTOGEN to set a parameter above a specified value. MAX_ refers to the maximum value to which a parameter can be set by AUTOGEN. The following example sets the maximum value to 400,000:

```
MAX_PAGEDYN = 400000
```

15.5.1.4 Specifying an Absolute Value

Use this method to specify a value for a parameter that AUTOGEN does not calculate. (Refer to the AUTOGEN section of the *OpenVMS System Management Utilities Reference Manual* for a table of the system parameters modified in AUTOGEN calculations.)

Note

Compaq strongly recommends that you use this method only for parameters that describe the system environment (for example, SCSNODE and SCSSYSTEMID). For the parameters that AUTOGEN calculates, specifying a value with this method disables AUTOGEN's calculations. Instead of specifying an absolute value, use one of the following methods:

- Specify a minimum value with the MIN_ prefix
 - Specify a maximum value with the MAX_ prefix
 - Increase the value with the ADD_ prefix
-

To specify an absolute parameter value, add an assignment statement in the following format to MODPARAMS.DAT:

```
parameter = parameter-value ! comment
```

Managing System Parameters

15.5 Modifying System Parameters with AUTOGEN

For example, the following command assigns the node name BIGVAX to the SCSNODE parameter:

```
SCSNODE = "BIGVAX" ! the node name
```

Note

String values must be within quotation marks ("").

15.5.1.5 Defining the Number of VAXcluster Nodes (VAX Only)

In a VAXcluster environment, use the NUM_NODES symbol to prevent temporary changes in VAXcluster membership from affecting AUTOGEN's calculation of VAXcluster-related parameter values. Define the NUM_NODES symbol in MODPARAMS.DAT to specify the number of nodes that are to run in the VAXcluster. AUTOGEN uses this value to set parameters that are affected by the number of VAXcluster nodes. For example, include the following line in MODPARAMS.DAT:

```
NUM_NODES = 30
```

15.5.1.6 Defining the Number of Ethernet Adapters (VAX Only)

In a VAXcluster environment, define the NUM_ETHERADAPT symbol in MODPARAMS.DAT to specify the total number of Ethernet adapters in the VAXcluster system. For example, include the following line in MODPARAMS.DAT:

```
NUM_ETHERADAPT = 40
```

15.5.1.7 Presetting Parameter Values Before Adding Memory (VAX Only)

On VAX systems, if you plan to upgrade your system hardware by adding a large amount (512 MB or more) of memory, you might want to preset your system parameters to values appropriate for the additional memory. Presetting your system parameters minimizes the possibility of memory upgrade problems caused by inappropriate parameter values.

How to Perform This Task

Perform the following steps:

1. Add a line to SYS\$SYSTEM:MODPARAMS.DAT in the following format:

```
MEMSIZE = total-number-of-pages-of-memory-after-upgrade
```

For example:

```
MEMSIZE = 2048 * 1024 ! (2048 page per MB * 1GB of memory)
```

2. Run AUTOGEN to the SETPARAMS phase.
3. Perform the hardware upgrade to add the additional memory.
4. Edit MODPARAMS.DAT to remove the line added in step 1.

15.5.1.8 Overriding Parameters Related to DECnet

To override AUTOGEN's observations regarding the presence (or absence) of DECnet, set the MODPARAMS.DAT parameter LOAD_DECNET_IMAGES to TRUE (or FALSE). Controlling the setting is useful for sites that have no synchronous network hardware but want to run asynchronous DECnet.

Managing System Parameters

15.5 Modifying System Parameters with AUTOGEN

15.5.1.9 Values Set for NPAGEDYN and NPAGEVIR

For the benefit of OpenVMS VAX systems with limited physical memory, AUTOGEN logs a warning message in its report if NPAGEDYN exceeds 10 percent of physical memory or if NPAGEVIR exceeds 33 percent of physical memory.

AUTOGEN also limits its own calculated value for NPAGEDYN to 20 percent of physical memory, and limits NPAGEVIR to 50 percent of physical memory. These calculated values are adequate for most workstations and systems with 16 or fewer megabytes of physical memory. If your system requires a larger value, you can override the AUTOGEN calculated values by setting higher values in MODPARAMS.DAT.

15.5.2 Specifying a Minimum Required Age for Feedback (VAX Only)

On VAX systems, AUTOGEN feedback is useful only when a system has been running long enough to accurately reflect the system's normal work load. By default, AUTOGEN uses feedback if the data is older than 24 hours. On VAX systems, you can define the logical name AGEN\$FEEDBACK_REQ_TIME to specify, in hours, a different minimum age required for feedback. AUTOGEN uses this value to determine whether the feedback is to be used.

For example, you might define the logical name as follows, to indicate that AUTOGEN should use feedback if it is older than 19 hours:

```
$ DEFINE/SYSTEM AGEN$FEEDBACK_REQ_TIME 19
```

To define this logical name each time the system starts up, add this command to SYLOGICALS.COM.

15.5.3 Including an External Parameter File in MODPARAMS.DAT

You can include external parameter files in MODPARAMS.DAT. For example, you might want to set a system parameter to the same value on all nodes in a VAXcluster or an OpenVMS Cluster environment; you might also want to specify node-specific values for other system parameters. You could specify the cluster-common values in a separate cluster-common file and include this cluster-common file in the MODPARAMS.DAT file on each system in the cluster.

To include a parameter file, place a command in the following format in MODPARAMS.DAT, or in any parameter file that is included in MODPARAMS.DAT:

```
AGEN$INCLUDE_PARAMS full-directory-spec:filename
```

Example

To include a cluster-common parameter file named CLUSTERPARAMS.DAT, create a common parameter file with the following name:

```
SYS$COMMON:[SYSEXE]CLUSTERPARAMS.DAT
```

Add the following line in the MODPARAMS.DAT file in the system-specific directory of each cluster:

```
AGEN$INCLUDE_PARAMS SYS$COMMON:[SYSEXE]CLUSTERPARAMS.DAT
```

15.5.4 Turning Off Logging of DCL Statements

The contents of MODPARAMS.DAT are evaluated as DCL statements; you can make assignments to symbols with names that are not system parameters (for example, scratch variables or conditional assignments based on other values). Traditionally, every such assignment is logged in AGEN\$PARAMS.REPORT, sometimes creating a large file with many logging statements that do not interest users.

You can prefix any assignments that you prefer not to log in AGEN\$PARAMS.REPORT with a dollar sign (\$). When AUTOGEN encounters a MODPARAMS.DAT record beginning with a dollar sign, it does not check the list of known system parameters and does not log this record to AGEN\$PARAMS.REPORT.

15.6 Automating AUTOGEN Reports

Compaq recommends that you create a batch-oriented command procedure to automatically run AUTOGEN on a regular basis and send the resulting feedback reports to an appropriate Mail account. Example 15–2 provides a sample command procedure.

Note

This command procedure runs AUTOGEN only to recommend system parameter values and send you a report. It does not run AUTOGEN to change system parameters or reboot the system. If, after reviewing the report, you decide to change system parameters, follow the instructions in Section 15.6.1.

The command procedure in Example 15–2 runs two passes of AUTOGEN. On the first pass, AUTOGEN runs during peak workload times to collect data on realistic system work loads. This pass does not degrade system performance. On the second pass, AUTOGEN runs during off-peak hours to interpret the data collected in the first stage.

The procedure sends the resulting report, contained in the file AGEN\$PARAMS.REPORT, to the SYSTEM account. Review this report on a regular basis to see whether the load on the system has changed.

Example 15–2 shows a sample command procedure. Use this procedure only as an example; create a similar command procedure as necessary to meet the needs of your configuration.

Example 15–2 Sample AUTOGEN Command Procedure

```
$ BEGIN$: ! ++++++++ AGEN_BATCH.COM ++++++++  
$ on warning then goto error$  
$ on control_y then goto error$
```

(continued on next page)

Managing System Parameters

15.6 Automating AUTOGEN Reports

Example 15–2 (Cont.) Sample AUTOGEN Command Procedure

```
$!  
$! Setup process  
$!  
$! Set process information  
$ set process/priv=all/name="AUTOGEN Batch"  
$! Keep log files to a reasonable amount  
$ purge/keep=5 AGEN_Batch.log  
$ time = f$time() ! Fetch current time  
$ hour = f$integer(f$cvtime(time,"hour")) ! Get hour  
$ today = f$cvtime(time,"WEEKDAY") ! Get Day of the week  
$ if f$integer(f$cvtime(time,"minute")) .ge. 30 then hour = hour + 1  
$!  
$! Start of working day...  
$!  
$ 1AM$:  
$ if hour .le. 2  
$ then  
$ next time = "today+0-14"  
$ gosub submit$ ! Resubmit yourself  
$ set noon  
$!  
$! Run AUTOGEN to TESTFILES using the parameter values collected earlier  
$! in the day (i.e., yesterday at 2:00pm)  
$ if today .eqs. "Tuesday" .OR. today .eqs. "Thursday" .OR. -  
today .eqs. "Saturday"  
$ then  
$ @sys$update:autozen GETDATA TESTFILES feedback ②  
$ mail/sub="AUTOGEN Feedback Report for system-name" -  
sys$system:agen$params.report system ③  
$ ! Clean up  
$ purge/keep=7 sys$system:agen$feedback.report ④  
$ purge/keep=7 sys$system:agen$feedback.dat  
$ purge/keep=7 sys$system:params.dat  
$ purge/keep=7 sys$system:autozen.par  
$ purge/keep=7 sys$system:setparams.dat  
$ purge/keep=7 sys$system:agen$addhistory.tmp  
$ purge/keep=7 sys$system:agen$addhistory.dat  
$ endif  
$ goto end$  
$ endif  
$!  
$ 2PM$:  
$ if hour .le. 15  
$ then  
$ next time = "today+0-17"  
$ gosub submit$  
$ if today .eqs. "Monday" .OR. today .eqs. "Wednesday" .OR. -  
today .eqs. "Friday"  
$ then  
$ @sys$update:autozen SAVPARAMS SAVPARAMS feedback ①  
$ endif  
$ goto end$  
$ endif  
$!  
$ 5PM$:  
$ if hour .le. 18  
$ then  
$ next time = "tomorrow+0-1"  
$ gosub submit$  
$ endif
```

(continued on next page)

Managing System Parameters 15.6 Automating AUTOGEN Reports

Example 15–2 (Cont.) Sample AUTOGEN Command Procedure

```
$!  
$! End of working day...  
$!  
$ END$:      ! ----- BATCH.COM -----  
$ exit  
$!++  
$! Subroutines  
$!--  
$!  
$ SUBMIT$:  
$ submit/name="AGEN_Batch"/restart/noprint - ⑤  
  /log=AGEN_batch.log -  
  /queue=sys$batch/after=""'next_time'" sys$system:AGEN_batch.com  
$ return  
$!++  
$! Error handler  
$!--  
$!  
$ ERROR$:  
$ mail/sub="AGEN_BATCH.COM - Procedure failed." _nl: system  
$ goto end$
```

The commands in this procedure perform the following tasks:

- ① Executes the first pass of AUTOGEN during peak workload times to collect data on realistic work loads. This command runs a very fast image so it does not degrade system response.
- ② Executes the second pass of AUTOGEN during off-peak hours to interpret the data collected in the first pass.
- ③ Mails the resulting report file named AGEN\$PARAMS.REPORT to the SYSTEM account.
- ④ Cleans up the files created.
- ⑤ Resubmits the command procedure.

15.6.1 Changing Parameter Values After Reviewing AUTOGEN Reports

If the command procedure report described in Section 15.6 shows AUTOGEN's calculations are different from the current values, correct the tuning by executing AUTOGEN with one of the two following commands:

- If the system can be shut down and rebooted immediately, execute the following command:

```
$ @SYSSUPDATE:AUTOGEN GETDATA REBOOT FEEDBACK
```
- If the system cannot be shut down and rebooted immediately, execute the following command to reset the system parameters:

```
$ @SYSSUPDATE:AUTOGEN GETDATA SETPARAMS FEEDBACK
```

The new parameters will take effect the next time the system boots.

Managing System Parameters

15.7 Managing System Parameters with SYSMAN

15.7 Managing System Parameters with SYSMAN

Note

Compaq recommends that you use AUTOGEN to modify system parameters. For more information, see Section 15.5. To view system parameters for a group of nodes or change parameters temporarily, use the System Management utility (SYSMAN).

SYSMAN provides the ability to inspect and modify system parameters for an entire cluster or for any group of nodes, rather than just one system. The PARAMETERS commands available in SYSMAN duplicate the parameter functions of the System Generation utility (SYSGEN).

You can use SYSMAN to manage system parameters as follows:

Task	For More Information
Show parameter values	Section 15.7.2
Modify current values in the parameter file	Section 15.7.3
Modify active values on a running system ¹	Section 15.7.4

¹Applies only to dynamic system parameters.

SYSMAN provides the commands and functions shown in Table 15–3.

Table 15–3 SYSMAN PARAMETERS Commands

Command	Function
PARAMETERS SHOW	Displays parameter values. Requires the name of the parameter.
PARAMETERS USE	Reads a set of parameters from memory or disk into the work area for inspection or modification. Requires a file name or the additional parameters ACTIVE or CURRENT.
PARAMETERS SET	Changes parameter values only in the work area; more permanent modification requires the PARAMETERS WRITE command. Requires the name and value of the parameter.
PARAMETERS WRITE	Writes the content of the work area to memory or to disk. Requires a file name or the additional parameters ACTIVE or CURRENT.

For more information about the temporary work area, see the next section.

15.7.1 Understanding Parameter Values and SYSMAN

It helps to understand the different system parameter values explained in Section 15.1.1. Briefly, **current values** are stored in the default parameter file on disk. **Active values** are stored in memory and are used while the system is running. In addition to these values, SYSMAN writes a temporary copy into its own work area on disk. Figure 15–2 illustrates these different sets of values and how SYSMAN commands affect them. In this figure:

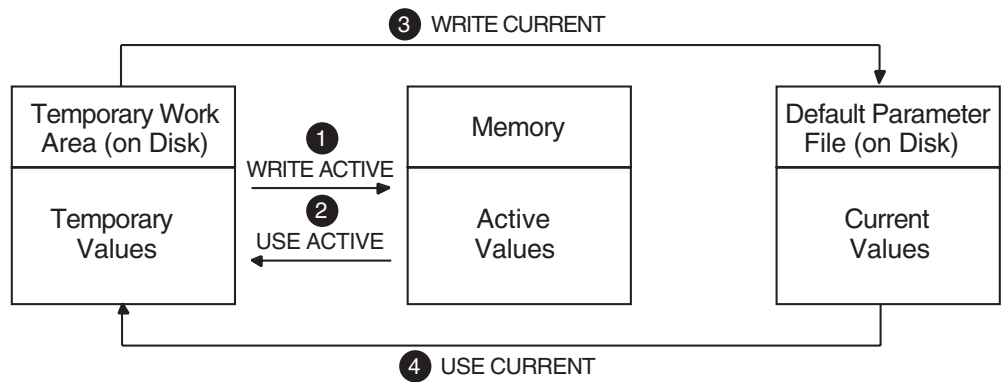
- 1 WRITE ACTIVE writes temporary parameter values to memory.
- 2 USE ACTIVE reads values from memory into the work area, where you can modify them.

Managing System Parameters

15.7 Managing System Parameters with SYSMAN

- ③ WRITE CURRENT writes temporary parameter values to disk, where they become current values. They become active the next time the system boots.
- ④ USE CURRENT reads the current values from disk into the work area, where you can modify them.

Figure 15–2 SYSMAN Temporary, Active, and Current Parameter Values



ZK-5275A-GE

During a typical session, you can display and change values in the following sequence:

1. Read values into SYSMAN's temporary work space with the USE command. USE ACTIVE reads in active values. USE CURRENT reads in current values.
2. Display the parameter values with the SHOW command.
3. Change a value with the SET command. You must use the WRITE command to activate the value.
4. Make the change effective with the WRITE command:
 - WRITE ACTIVE writes the value to the set of active values. (You can change an active value only if the parameter is a dynamic parameter.)
 - WRITE CURRENT writes the value to the set of current values.

For a list of all the system parameters, refer to the *OpenVMS System Management Utilities Reference Manual*.

15.7.2 Showing Parameter Values with SYSMAN

Use the SYSMAN command PARAMETERS SHOW to display parameter values for all the nodes in a cluster.

Examples

1. The following example shows one method to display information about parameters. In this case, using the /LGI qualifier displays all login security control parameters. You can display many categories of parameters, such as /ACP, /ALL, and /SPECIAL. Refer to the *OpenVMS System Management Utilities Reference Manual* for a complete list of parameters and parameter categories.

Managing System Parameters

15.7 Managing System Parameters with SYSMAN

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS SHOW/LGI
```

```
Parameters in use: Active
Parameter Name      Current  Default  Min.    Max.    Unit  Dynamic
-----
LGI_BRK_TERM        0         1         0         1 Boolean D
LGI_BRK_DISUSER     0         0         0         1 Boolean D
LGI_PWD_TMO         30        30         0        255 Seconds D
LGI_RETRY_LIM       3         3         0        255 Tries D
LGI_RETRY_TMO       20        20         0        255 Seconds D
LGI_BRK_LIM         5         5         0        255 Failures D
LGI_BRK_TMO         300       300         0         -1 Seconds D
LGI_HID_TIM         300       300         0         -1 Seconds D
```

- The following example invokes SYSMAN and specifies the environment to be the local cluster, which consists of NODE21 and NODE22. The example also displays the active value for the LGI_BRK_TMO parameter, which controls the number of seconds that a user, terminal, or node is permitted to attempt login. In this case, it is 600.

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
%SYSMAN-I-ENV, Current command environment:
Clusterwide on local cluster
Username MORIN will be used on nonlocal nodes
SYSMAN> PARAMETERS SHOW LGI_BRK_TMO
```

```
Node NODE21: Parameters in use: ACTIVE
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO         600       300         0         -1 Seconds D

Node NODE22: Parameters in use: ACTIVE
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO         600       300         0         -1 Seconds D
```

15.7.3 Modifying a Parameter File with SYSMAN

Use the SYSMAN command PARAMETERS WRITE to write system parameter values and the name of the site-independent startup command procedure to your choice of parameter file or the current system parameter file on disk.

The PARAMETERS WRITE CURRENT command sends a message to OPCOM to record the event, unless you have changed the system message format with the DCL command SET MESSAGE.

Note

The PARAMETERS WRITE CURRENT command writes *all* of the active or current parameter values—not just the one you may be working on—to disk.

Examples

- The following example creates a new parameter specification file:

```
SYSMAN> PARAMETERS WRITE SYS$SYSTEM:NEWPARAM
```

- When used with the PARAMETERS SET command, the PARAMETERS WRITE command modifies the current system parameter file on disk:

```
SYSMAN> PARAMETERS SET LGI_BRK_TMO 300
SYSMAN> PARAMETERS WRITE CURRENT
```

Managing System Parameters

15.7 Managing System Parameters with SYSMAN

15.7.4 Modifying Active Values with SYSMAN

Using the SYSMAN commands PARAMETERS SET, PARAMETERS WRITE, and PARAMETERS USE enables you to modify active parameter values.

Modifying active values immediately affects dynamic parameters by changing their values in memory. Appendix C of the *OpenVMS System Management Utilities Reference Manual* identifies dynamic parameters, as does the SYSMAN command PARAMETERS SHOW/DYNAMIC. Values for nondynamic parameters cannot be changed while the system is running.

Modifying active values does not affect current values in the system parameter file on disk, because the next time you boot the system, the values on disk are established as the active values.

If you set new active parameter values and you want to use the new values for subsequent boot operations, write the new values to the current parameter file with the PARAMETERS WRITE CURRENT command, as shown in the Examples section.

Caution

Parameter values modified with SYSMAN will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with SYSMAN, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 15.5.1 to specify the new parameter values.

Examples

1. The following example changes the LGI_BRK_TMO value to 300 in the work area, writes this change into memory as an active value, and displays the active value:

```
SYSMAN> PARAMETERS SET LGI_BRK_TMO 300
SYSMAN> PARAMETERS WRITE ACTIVE
SYSMAN> PARAMETERS SHOW LGI_BRK_TMO

Node NODE21: Parameters in use: ACTIVE
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO         300     300      0        -1 Seconds  D

Node NODE22: Parameters in use: ACTIVE
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO         300     300      0        -1 Seconds  D
```

2. The following example calls the current parameter values, including LGI_BRK_TMO, from disk to the work area, then displays LGI_BRK_TMO. In this example, the current value on disk is 600.

```
SYSMAN> PARAMETERS USE CURRENT
SYSMAN> PARAMETERS SHOW LGI_BRK_TMO

Node NODE21: Parameters in use: CURRENT
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO         600     300      0        -1 Seconds  D
```

Managing System Parameters

15.7 Managing System Parameters with SYSMAN

```

Node NODE22: Parameters in use: CURRENT
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO        600     300      0        -1 Seconds  D

```

- The next example writes the LGI_BRK_TMO value of 600 from the work area to memory, where it becomes the active value on the running system. Note that the command `PARAMETER WRITE ACTIVE` writes all the parameter values from the work area into memory, not just the value of LGI_BRK_TMO.

```

SYSMAN> PARAMETERS WRITE ACTIVE
SYSMAN> PARAMETERS USE ACTIVE
SYSMAN> PARAMETERS SHOW LGI_BRK_TMO

```

```

Node NODE21: Parameters in use: ACTIVE
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO        600     300      0        -1 Seconds  D

Node NODE22: Parameters in use: ACTIVE
Parameter Name      Current  Default  Minimum  Maximum Unit  Dynamic
-----
LGI_BRK_TMO        600     300      0        -1 Seconds  D

```

15.8 Managing System Parameters with SYSGEN

Note

Compaq recommends that you use `AUTOGEN` to modify system parameters. For more information, see Section 15.5. If for some reason you cannot use `AUTOGEN`, Compaq recommends that you use `SYSMAN`. For more information, see Section 15.7.

Although it is not the recommended method, you can also use the System Generation utility (`SYSGEN`) to manage system parameters as follows:

Task	For More Information
Show parameter values	Section 15.8.2
Modify current values in the default parameter file	Section 15.8.3
Modify active values on a running system ¹	Section 15.8.4
Create a new parameter file	Section 15.8.5

¹Applies only to the dynamic system parameters.

`SYSGEN` provides the commands shown in Table 15–4 for managing system parameters. Refer to the `SYSGEN` section of the *OpenVMS System Management Utilities Reference Manual* for detailed descriptions of `SYSGEN` commands.

Managing System Parameters

15.8 Managing System Parameters with SYSGEN

Table 15–4 SYSGEN Commands Used with System Parameters

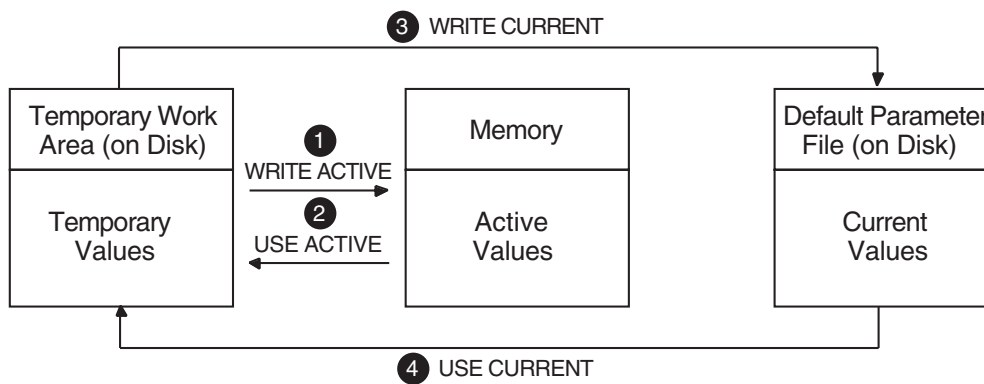
Command	Function
SHOW	Displays parameter values.
USE	Reads a set of values from memory or disk into a temporary work area for inspection or modification.
SET	Changes parameter values only in the work area; more permanent modification requires the WRITE command.
WRITE	Writes the content of the work area to memory or to disk.

For more information about the temporary work area, see the next section.

15.8.1 Understanding Parameter Values and SYSGEN

You should understand the different system parameter values explained in Section 15.1.1. Briefly, **current values** are stored in the default parameter file on disk. **Active values** are stored in memory and are used while the system is running. In addition to these values, SYSGEN writes a temporary copy into its own work area on disk. Figure 15–3 illustrates these different sets of values and shows how SYSGEN commands affect them.

Figure 15–3 SYSGEN Temporary, Active, and Current Parameter Values



In a typical session, you might display and change values in the following sequence:

1. Read values into SYSGEN's temporary work space with the USE command. USE ACTIVE reads in active values. USE CURRENT reads in current values.
2. Display the parameter values with the SHOW command.
3. Change a value with the SET command. (Note, however that the SET command only changes the value in SYSGEN's temporary work area.)
4. Make the change effective with the WRITE command:
 - WRITE ACTIVE writes the value to the set of active values in memory. (You can change an active value only if the parameter is a dynamic parameter.)
 - WRITE CURRENT writes the value to the set of current values on disk.

Managing System Parameters

15.8 Managing System Parameters with SYSGEN

For a list of all the system parameters, refer to the *OpenVMS System Management Utilities Reference Manual*.

15.8.2 Showing Parameter Values with SYSGEN

To display values for system parameters, perform the following steps:

1. Invoke SYSGEN by entering the following command:

```
$ RUN SYS$SYSTEM:SYSGEN
```

2. Enter the USE command to specify which values you want to display, as follows:

To Display	Enter
Active values	USE ACTIVE
Current values	USE CURRENT
Values from another parameter file	USE <i>file-spec</i>

For *file-spec*, specify the parameter file from which you want to display values; for example, USE SYS\$SYSTEM:ALTPARAMS.DAT

3. Enter a SHOW command in the following format:

```
SHOW [/qualifier] [parameter-name]
```

Specify qualifiers to display parameters grouped by type. For example:

To Display Values For	Enter
The WSMAX parameter	SHOW WSMAX
All dynamic parameters	SHOW/DYNAMIC
All parameters in the TTY category	SHOW/TTY
All parameters	SHOW/ALL

For more information about the SYSGEN SHOW command and qualifiers, refer to the SYSGEN section of the *OpenVMS System Management Utilities Reference Manual*.

Example

The following example uses SYSGEN to show the current values of all TTY system parameters:

Managing System Parameters

15.8 Managing System Parameters with SYSGEN

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> USE CURRENT
SYSGEN> SHOW/TTY
```

```
Parameters in use: Current1
Parameter Name      Current      Default      Min.      Max. Unit      Dynamic
-----
2
TTY_SCANDELTA       10000000    10000000    100000    -1 100Ns
TTY_DIALTYPE        0            0            0          255 Bit-Encode
TTY_SPEED           15           15           1          16 Special
TTY_RSPEED          0            0            0          16 Special
TTY_PARITY          24           24           0          255 Special
TTY_BUF             80           80           0          65535 Characters
TTY_DEFCHAR         402657952   402657952   0          -1 Bit-Encode
TTY_DEFCHAR2        135178      4098         0          -1 Bit-Encode
TTY_TYPAHDSZ        78           78           0          -1 Bytes
TTY_ALTYPAMD        2048        200          0          32767 Bytes
TTY_ALTALARM        750         64           0          -1 Bytes
TTY_DMASIZE         64           64           0          -1 Bytes      D 8
TTY_CLASSNAME       "TTY"        "TTY"        "AA"       "ZZ" Ascii
TTY_SILOTIME        8            8            0          255 Ms
TTY_TIMEOUT         3600         900          0          -1 Seconds   D
TTY_AUTOCHAR        7            7            0          255 Character D
SYSGEN>
```

SYSGEN displays the following information:

- ¹ The values in use (in this example, current values)
- ² The name of the system parameter
- ³ The value requested (in this example, the current value). The heading of this column is always “Current,” regardless of whether it displays the current or active value of the parameter. In this context, “Current” refers to the value of this parameter *currently* in use, as specified by the USE command; it does not refer to the *current value* of the parameter stored on disk with the WRITE CURRENT command.
- ⁴ The default value
- ⁵ The minimum value
- ⁶ The maximum value
- ⁷ The unit of allocation
- ⁸ A “D,” if the system parameter is dynamic

15.8.3 Modifying the System Parameter File with SYSGEN

Caution

Parameter values modified with the System Generation utility (SYSGEN) will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with SYSGEN, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 15.5.1 to specify the new parameter values.

Managing System Parameters

15.8 Managing System Parameters with SYSGEN

Note

Although you can modify system parameter values with SYSGEN, Compaq recommends that you use AUTOGEN. For more information, see Section 15.5.

If you cannot use AUTOGEN, Compaq recommends that you use the System Management utility (SYSMAN) to modify system parameters. For more information, see Section 15.7.

Modifying the current values in the default system parameter file has no immediate effect on active values on a running system. However, during subsequent boot operations, the system is initialized with the new values.

Example

The following example modifies the TTY_TIMEOUT parameter value in the VAX system parameter file:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> USE CURRENT
SYSGEN> SET TTY_TIMEOUT 3600
SYSGEN> WRITE CURRENT
%OPCOM, 15-APR-2000 16:04:06.30, message from user SYSTEM
%SYSGEN-I-WRITECUR, CURRENT system parameters modified by process
ID 00160030 into file VAXVMS$SYS.PAR
SYSGEN> EXIT
```

15.8.4 Modifying Active Values with SYSGEN

Caution

Parameter values modified with SYSGEN will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with SYSGEN, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 15.5.1 to specify the new parameter value.

Note

Although you can modify system parameter values with SYSGEN, Compaq recommends that you use AUTOGEN or the System Management utility (SYSMAN). For more information, see Section 15.7.

Modifying active values immediately affects dynamic parameters by changing their values in memory. Appendix C of the *OpenVMS System Management Utilities Reference Manual* identifies dynamic parameters (as does the SYSGEN command SHOW/DYNAMIC). You cannot change values for nondynamic parameters while the system is running.

Modifying active values does not affect the current values in the system parameter file on disk. The next time you boot the system, the old current values are established as the active values.

Managing System Parameters

15.8 Managing System Parameters with SYSGEN

If you set new active parameter values (by entering WRITE ACTIVE) and you want to use the new values for subsequent boot operations, you must write the new values to the current parameter file on disk by entering the WRITE CURRENT command, as explained in Section 15.8.3. If the parameters are not dynamic parameters, you must enter the WRITE CURRENT command and reboot the system.

When you change active parameters with SYSGEN, the operator communication manager (OPCOM) writes a message to the operator log and the operator console, unless you have changed the system message format with the DCL command SET MESSAGE.

Examples

1. The following example modifies the active value of the PFCDEFAULT parameter:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> SET PFCDEFAULT 127
SYSGEN> WRITE ACTIVE
%OPCOM, 15-APR-2000 16:04:06.30, message from user SYSTEM
%SYSGEN-I-WRITEACT, ACTIVE system parameters modified by process
ID 00160030
SYSGEN> EXIT
```

2. The following example modifies the active value of the PFCDEFAULT parameter and also writes it to the Alpha system parameter file, so it will be used when the system reboots:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> SET PFCDEFAULT 127
SYSGEN> WRITE ACTIVE
%OPCOM, 15-APR-2000 16:04:06.30, message from user SYSTEM
%SYSGEN-I-WRITEACT, ACTIVE system parameters modified by process
ID 00160030
SYSGEN> WRITE CURRENT
%OPCOM, 15-APR-2000 16:04:06.30, message from user SYSTEM
%SYSGEN-I-WRITECUR, CURRENT system parameters modified by process
ID 00160030 into file ALPHAVMSSYS.PAR
SYSGEN> EXIT
```

15.8.5 Creating a New Parameter File with SYSGEN

Creating a new parameter file has no effect on the running system. During a subsequent conversational boot operation, however, you can initialize the active system with the values of the new file.

How to Perform This Task

1. Invoke SYSGEN by entering the following commands:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
```

2. Enter a command in the following format to write a copy of a parameter file into SYSGEN's temporary workspace:

```
USE file-spec
```

where *file-spec* is the file specification for the parameter file to be used as a base. Modify the values in this file to create a new parameter file.

Managing System Parameters

15.8 Managing System Parameters with SYSGEN

3. Enter commands in the following form to modify values as needed:
SET parameter-name value
where *parameter-name* specifies the name of the parameter to be changed, and *value* specifies the new value for the parameter.
4. Specify a command in the following format to write the values to a new parameter file:
WRITE file-spec
where *file-spec* is the file specification for the parameter file to be created.
5. Exit SYSGEN.

Caution

Parameter values modified with SYSGEN are overridden by the AUTOGEN command procedure. To keep parameter modifications made with SYSGEN, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 15.5.1 to specify the new parameter values.

Examples

1. The following example creates a new version of the parameter file PARAMS.PAR with a new value for the TTY_TIMEOUT parameter:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> USE SYS$MANAGER:PARAMS.PAR
SYSGEN> SET TTY_TIMEOUT 3600
SYSGEN> WRITE SYS$MANAGER:PARAMS.PAR
SYSGEN> EXIT
```
2. The following example creates a file named SYS\$SYSTEM:OURSITE.PAR, using the PARAMS.PAR file as a base:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYSGEN
SYSGEN> USE SYS$MANAGER:PARAMS.PAR
SYSGEN> SET TTY_TIMEOUT 1000
SYSGEN> WRITE OURSITE.PAR
SYSGEN> EXIT
```

15.9 Modifying System Parameters with a Conversational Boot

Note

Although you can modify system parameters with a conversational boot, Compaq recommends that you use AUTOGEN or the System Management utility (SYSMAN). For more information, see Section 15.5 and Section 15.7.

Use a conversational boot only to change isolated system parameters *temporarily* or in an emergency. For example, during a system upgrade, you would use a conversational boot to modify STARTUP_P1 to use a minimum startup.

Managing System Parameters

15.9 Modifying System Parameters with a Conversational Boot

Remember that if you change a value and do not add the changed value to the AUTOGEN parameter file MODPARAMS.DAT, AUTOGEN will overwrite the value the next time AUTOGEN executes.

With a conversational boot operation, you can modify the active parameter values in the following ways before the system boots:

Task	For More Information
Modify active values for individual parameters	Section 4.2.1
Initialize active values using values stored in a parameter file other than the default parameter file	Section 4.2.2
Reinitialize active values using default values	Section 4.4.1

At the end of the conversational boot, the default system parameter file is modified to store the new active parameter values.

Caution

Parameter values modified with a conversational boot will be overridden by the AUTOGEN command procedure. To keep parameter modifications made with a conversational boot, edit the file SYS\$SYSTEM:MODPARAMS.DAT as explained in Section 15.5.1 to specify the new parameter values.

15.10 Tuning BAP System Parameters

OpenVMS Alpha Version 7.1 and later contains system parameters that control the operation of bus-addressable pool (BAP).

The CIPCA, CIXCD, KFMSB, and Qlogic 1020ISP adapters are some of the adapters that use bus-addressable pool to improve performance. BAP is a non-paged dynamic, physical-address-filtered memory pool used to overcome I/O bus and 32-bit adapter physical addressing limits.

The following table lists the system parameters that control BAP operation along with their default values:

System Parameter	Default Value
NPAG_BAP_MIN	0
NPAG_BAP_MAX	0
NPAG_BAP_MIN_PA	0
NPAG_BAP_MAX_PA	-1

The default values of these parameters allow the system to boot with any configuration. When AUTOGEN is run on a configured system, it resets these parameters to values that should enhance performance for the current system configuration.

Managing System Parameters

15.10 Tuning BAP System Parameters

If the system fails to boot after an installation, upgrade, or configuration change, and displays a message that refers to incorrect BAP parameters, Compaq recommends that you perform the following steps:

1. Reset the BAP parameters to the default values.
2. Reboot the system.
3. Allow the installation procedure to run AUTOGEN, or manually run AUTOGEN yourself.

A typical AUTOGEN with FEEDBACK command to set these parameters follows:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS SETPARAMS FEEDBACK
```

Note

These parameters are critical. Compaq recommends that you run AUTOGEN as described to ensure that they are set correctly.

If you prefer not to use this command because you want to adjust only the BAP parameters settings, use the following procedure:

1. Boot the system using the default BAP parameter values.
2. Manually run SYS\$SYSTEM:AGEN\$FEEDBACK.EXE:


```
$ RUN SYS$SYSTEM:AGEN$FEEDBACK.EXE
```
3. Search SYS\$SYSTEM:AGEN\$FEEDBACK.DAT for the BAP_* system parameter values:


```
$ SEARCH SYS$SYSTEM:AGEN$FEEDBACK.DAT "BAP_"
```
4. Run SYSGEN to set the following system parameters with the BAP values you obtained in Step 3:

AGEN\$FEEDBACK Data	System Parameter	Units
BAP_MIN	NPAG_BAP_MIN	bytes
BAP_MAX	NPAG_BAP_MAX	bytes
BAP_MIN_PA	NPAG_BAP_MIN_PA	Mbytes ¹
BAP_MAX_PA	NPAG_BAP_MAX_PA	Mbytes ¹

¹On OpenVMS Alpha systems prior to Version 7.2, the value of this parameter is specified in bytes.

The BAP allocation amount (specified by BAP_MIN and BAP_MAX) depends on the adapter type, the number of adapters, and the version of the operating system. The physical address range (specified by BAP_MIN_PA and BAP_MAX_PA) depends on the adapter type and the way the Galaxy logical partitions, if any, are defined.

Note

If you manually set parameters NPAG_BAP_MIN_PA and NPAG_BAP_MAX_PA, be sure to specify the value for each parameter in the correct units (bytes or megabytes) for your operating system version.

Managing Page, Swap, and Dump Files

Page, swap, and dump files are created by default. However, you should understand these files. In addition, you might want to change them to meet the needs of your site.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Displaying information about page and swap files	Section 16.3
Calculating appropriate sizes for files	Section 16.4
Minimizing dump file size when disk space is insufficient	Section 16.5
Writing a dump file on a system disk with multiple paths or shadow set members	Section 16.6
Writing the dump file to a device other than the system disk	Section 16.7
Using SDA to analyze the contents of a crash dump	Section 16.8
‡Using SDA CLUE commands to obtain and analyze summary crash dump information	Section 16.9
†Using CLUE to obtain historical information about crash dumps	Section 16.10
Saving the contents of the system dump file after a system failure	Section 16.11
Copying dump files to tape or disk	Section 16.12
Freeing dump information from the page file	Section 16.13
Installing page and swap files	Section 16.14
Removing page, swap, and dump files	Section 16.15
Creating and modifying page, swap, and dump files	Section 16.16
Controlling access to process dumps	Section 16.17.2
†VAX specific	
‡Alpha specific	

This chapter explains the following concepts:

Concept	Section
Understanding dump files	Section 16.1
Understanding page and swap files	Section 16.2
Understanding the order of information in a selective system dump	Section 16.5.1
‡Understanding SDA CLUE	Section 16.9.1
‡Alpha specific	

Managing Page, Swap, and Dump Files

Concept	Section
†Understanding CLUE	Section 16.10.1
Understanding process dumps	Section 16.17
†VAX specific	

16.1 Understanding Dump Files

When the operating system detects an unrecoverable error or an inconsistency within itself that causes the system to fail, it writes the contents of the error log buffers, processor registers, and memory into the **system dump file**, overwriting its previous contents.

The contents of error log buffers are also written to the **error log dump file**. The error log dump file is provided so that the system can be updated on reboot to include error log entries that were created but not written at the time of a system crash.

System Dump File

When writing the system dump file, the system displays console messages and information about the error or inconsistency. The last message tells you that the dump file was successfully written.

Caution

Be sure to wait until you see the termination message before using the console terminal to halt the system. If you do not wait, your system might not save a complete system dump file.

Console messages and the system dump file are important sources of information in determining the cause of a system failure. Use the contents in the following ways:

- Use the System Dump Analyzer utility (SDA) to analyze the contents of the dump and determine the cause of a failure.
- On Alpha systems, use SDA CLUE commands to obtain and analyze summary dump file information.
- On VAX systems, use CLUE to obtain historical information from system dump files.
- Contact your Compaq support representative, indicating that you have made a copy of the contents of the system dump.

The default system dump file, SYS\$SPECIFIC:[SYSEXE]SYSDUMP.DMP, is created during installation. (You do not need a system dump file to run the operating system. However, you must have a system dump file to diagnose system crashes.) AUTOGEN automatically determines an appropriate size for the system dump file for your hardware configuration and system parameters. Refer to Section 16.5 for information about minimizing system dump file size if disk space is insufficient.

Managing Page, Swap, and Dump Files

16.1 Understanding Dump Files

For special configurations or varying work loads, you can change the size of the system dump file. For information, see Section 16.16.1. You can write the system dump file on a disk other than the system disk. This is referred to as **dump off system disk** (DOSD). For more information, see Section 16.7.

Error Log Dump File

AUTOGEN creates the error log dump file during installation; its size depends on your configuration and system parameters. Error log dump files on VAX and Alpha systems have the following differences:

- On Alpha systems, the error log dump file is called SYS\$ERRLOG.DMP; it is located on the system disk. When an operator initiates a shutdown, the system writes the contents of error log buffers to the error log dump file and not to the system dump file. Thus, the last system crash dump is not overwritten.
- On VAX systems, the error log dump file is called SYSDUMP.DMP. How the system handles this file depends on whether you are using dump off system disk (DOSD):
 - With DOSD, error logs are written to a stub error log dump file, SYSDUMP.DMP, on the system disk. Also, both the error logs and system memory are written to the SYSDUMP.DMP file on the DOSD disk.
 - Without DOSD, both the error log and system memory are written to the SYSDUMP.DMP file on the system disk.

On VAX systems with or without DOSD, the last system crash dump is always overwritten when an operator initiates a shutdown.

16.1.1 Using the Page File to Store System Crash Dumps

The operating system uses the latest version of SYS\$SYSTEM:SYSDUMP.DMP to store system crash dumps. If SYSDUMP.DMP does not exist in SYS\$SYSTEM, the operating system uses the system **page file**, SYS\$SYSTEM:PAGEFILE.SYS, overwriting the contents of that file.

If the SAVEDUMP system parameter is set, the crash dump is retained in PAGEFILE.SYS when the system is booted. If SAVEDUMP is clear, the system uses the page file for paging; any dump written to the page file is lost.

If you use SYS\$SYSTEM:PAGEFILE.SYS to capture system crash dumps, you should later free the space occupied by the dump for use in system paging, with either of the following methods:

- Use the SDA COPY command to copy the dump from the page file to a different file.
- Use the SDA RELEASE command to delete the information from the page file.

For detailed instructions, see Section 16.13.

Include the appropriate commands in the SYSTARTUP_VMS.COM startup command procedure to free dump information from the page file each time the system reboots.

Managing Page, Swap, and Dump Files

16.1 Understanding Dump Files

Caution

Be careful when using the page file for selective dumps. Selective dumps use up all available space. If your page file is small, selective dump information might fill the entire page file, leaving no space for paging during system boot. This can cause the system to hang during reboot.

16.1.2 Understanding Types of System Dumps

The two types of system dumps are physical and selective. Table 16–1 defines physical and selective system dumps. Table 16–3 compares the information available in physical and selective system dump files.

Table 16–1 Definitions of Physical and Selective System Dumps

Type	Description
Physical dump	Writes the entire contents of physical memory to the system dump file. To ensure a useful physical dump, the system dump file must be large enough to contain all of physical memory.
Selective dump	Stores those portions of memory most likely to be useful in crash dump analysis. A selective system dump is useful when disk space is not available to hold all of physical memory.

Requirements for Creating a Useful System Dump

The following requirements must be met for the operating system to write a useful system dump file:

- The system parameter DUMPBUG must be set to 1 (the default value).
- If the system parameter SAVEDUMP is set to 0 (the default) the file `SYS$SPECIFIC:[SYSEXE]SYSDUMP.DMP` must exist on the system disk.
- If the file `SYS$SPECIFIC:[SYSEXE]SYSDUMP.DMP` does not exist on the system disk, the page file must be used to store the dump. The system parameter SAVEDUMP must be set to 1 and the file `SYS$SPECIFIC:[SYSEXE]PAGEFILE.SYS` must exist on the system disk.
- If the file `SYS$SPECIFIC:[SYSEXE]SYSDUMP.DMP` does not exist and `SYS$SPECIFIC:[SYSEXE]PAGEFILE.SYS` is not available for system dumps, you must create the dump file on an alternate disk (see Section 16.7).
- If sufficient disk space is not available to allow a system dump file that can hold all of memory, the system parameter DUMPSTYLE must be set to the appropriate value to store a selective system dump. For more information, refer to Section 16.5 and the *OpenVMS System Management Utilities Reference Manual*.
- The system dump file (or page file if the SAVEDUMP system parameter is set) must be large enough to hold all information that is to be written if the system fails.

If the system parameter DUMPBUG is set, AUTOGEN automatically sizes SYSDUMP.DMP if enough disk space is available.

If the system parameter SAVEDUMP is set, AUTOGEN performs no operations on the system dump file.

Managing Page, Swap, and Dump Files

16.1 Understanding Dump Files

AUTOGEN sizes the page file only for paging use, regardless of whether the SAVEDUMP system parameter is set.

BACKUP Considerations

A system dump file has the NOBACKUP attribute; therefore, the Backup utility (BACKUP) does not copy the file unless you use the qualifier /IGNORE=NOBACKUP when invoking BACKUP. When you use the SDA COPY command to copy the system dump file to another file, the operating system does not automatically set the new file to NOBACKUP. If you want to set the NOBACKUP attribute on the copy, use the SET FILE command with the /NOBACKUP qualifier as described in the *OpenVMS DCL Dictionary*.

Security Considerations

By default, SYS\$SYSTEM:SYSDUMP.DMP is protected against world access. Because a system dump file can contain privileged information, you should keep this level of protection on system dump files. Similarly, when you copy system dump files using the System Dump Analyzer utility (SDA) as explained in Section 16.11 and Section 16.13, be sure to protect the copy from world read access. For more information about file protection, refer to the *OpenVMS Guide to System Security*.

16.2 Understanding Page and Swap Files

As part of memory management, the operating system makes efficient use of physical memory by moving information between physical memory and files stored on disk. The system does this in two ways: paging and swapping. Table 16–2 defines these and related terms.

Table 16–2 Paging and Swapping Terminology

Term	Definition
Paging	A memory management operation that provides the efficient use of physical memory allotted to a <i>process</i> . Paging moves infrequently used portions of a process workspace out of physical memory to a file. For more information about paging, refer to the <i>OpenVMS Performance Management</i> manual.
Page file	The file to which the system writes paged portions of memory. The OpenVMS installation process creates a page file named SYS\$SYSTEM:PAGEFILE.SYS. If necessary, you can use SYS\$SYSTEM:PAGEFILE.SYS in place of the system crash dump file. For more information, see Section 16.1.1.
Swapping	A memory management operation that provides the efficient use of physical memory available for the <i>entire system</i> . Swapping moves the entire workspace of a less active process out of physical memory to a file. For more information about swapping, refer to the <i>OpenVMS Performance Management</i> manual.
Swap file	The file to which the system writes swapped portions of memory. The OpenVMS installation procedure creates a swap file named SYS\$SYSTEM:SWAPFILE.SYS.
Primary page and swap files	The default page and swap files created during OpenVMS installation. These files are named SYS\$SYSTEM:PAGEFILE.SYS and SYS\$SYSTEM:SWAPFILE.SYS.

(continued on next page)

Managing Page, Swap, and Dump Files

16.2 Understanding Page and Swap Files

Table 16–2 (Cont.) Paging and Swapping Terminology

Term	Definition
Secondary page and swap files	Additional page and swap files that you might create for performance or disk space reasons. If you kept the primary page and swap file on the system disk, the system uses the space in the secondary files for paging and swapping in addition to the space in the primary page and swap files. For information about creating secondary page and swap files, see Section 16.16.

Installing Files

Page and swap files must be installed before the system can use them. The system automatically installs the latest versions of `SYSD$SYSTEM:PAGEFILE.SYS` and `SWAPFILE.SYS` during startup. If you create secondary page and swap files, you must make sure the system installs them during startup. For more information about installing page and swap files, see Section 16.14.

File Sizes and Locations

AUTOGEN automatically determines appropriate sizes for the files for your hardware configuration and system parameters. For special configurations or varying work loads, you might want to change the size of the page or swap file. For information, see Section 16.16.1.

If your system does not require the page file for storing system crash dumps, you can move it off the system disk. However, you should keep one page file on the system disk, if possible, so that you can boot the system if another disk holding the page files becomes unavailable. The swap file can also be moved off the system disk.

16.3 Displaying Information About Page and Swap Files

The DCL command `SHOW MEMORY/FILES` displays information about the page and swap files existing on your system, including file names, sizes, and the amount of space used. For example:

```
$ SHOW MEMORY/FILES
      System Memory Resources on 19-JAN-2001 13:35:26.58

Swap File Usage (8KB pages):
  DISK$PAGE_DUMPS:[SYS0.SYSEXE]SWAPFILE.SYS;2
                                     Index      Free      Size
                                     1          7992     8248

Paging File Usage (8KB pages):
  DISK$PAGE_DUMPS:[SYS0.SYSEXE]PAGEFILE.SYS;1
                                     Index      Free      Size
                                     254       13722    16496

Total committed paging file usage:
                                     4870
```

The total committed paging file usage is the number of pages in the system that require pagefile space for paging. It can be bigger than the total number of pagefile pages that are available because it is unlikely that all the required space will be used for paging at one time.

16.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

16.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

When you install or upgrade the operating system, AUTOGEN automatically calculates appropriate sizes for your system; these sizes are based on your hardware configuration and your system parameters. However, you can manually calculate the sizes for these files. The following sections describe how to determine appropriate sizes for the system page, swap, and dump files.

16.4.1 Calculating System Dump File Size

Sufficient space in the system dump file is critical to saving a complete crash dump. The AUTOGEN command procedure calculates an appropriate size for your system dump file. However, if you want to manually calculate the system dump file size, use the following formula, which calculates the file size required to hold a physical dump.

For SYSDUMP.DMP

On VAX systems, use the following formula:

```
size-in-blocks(SYS$SYSTEM:SYSDUMP.DMP)
= size-in-pages(physical-memory)
+ number-of-error-log-buffers * blocks-per-buffer
+ 1
```

On Alpha systems, use the following formula:

```
size-in-blocks(SYS$SYSTEM:SYSDUMP.DMP)
= size-in-pages(physical-memory) * blocks-per-page
+ number-of-error-log-buffers * blocks-per-buffer
+ 10
```

where:

size-in-pages

Is the size of physical memory, in pages. Use the DCL command SHOW MEMORY to determine the total size of physical memory on your system.

blocks-per-page

Is the number of blocks per page of memory.

On Alpha systems, calculate the number of blocks per page of memory by dividing the system's page size by 512 (the size of a block). Use the following commands:

```
$ PAGESIZE==F$GETSYI ("PAGE SIZE")
$ BLOCKSPERPAGE=PAGESIZE/512
$ SHOW SYMBOL BLOCKSPERPAGE
```

number-of-error-log-buffers

Is the value of the system parameter ERRORLOGBUFFERS. This parameter sets the number of error log buffers to permanently allocate in memory.

blocks-per-buffer

Is the value of the system parameter ERLBUFFERPAGES. This parameter sets the number of pagelets (blocks) of memory in each buffer.

A large memory system or a system with small disk capacity might not be able to supply enough disk space for a full memory dump. Under these circumstances, you should set the system parameter DUMPSTYLE to the appropriate value to indicate that the system is to dump only selective information. For more information, see Section 16.5.

Managing Page, Swap, and Dump Files

16.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

For PAGEFILE.SYS

If SYS\$SYSTEM:SYSDUMP.DMP does not exist, the system writes crash dumps to the primary page file SYS\$SYSTEM:PAGEFILE.SYS. The AUTOGEN command procedure calculates an appropriate size for your page file. However, to manually calculate the minimum page file size required to hold crash dumps, use the following formula:

On VAX systems:

```
size-in-blocks(SYS$SYSTEM:PAGEFILE.SYS)
= size-in-pages(physical-memory)
+ number-of-error-log-buffers * blocks-per-buffer
+ 1
+ 1000
```

On Alpha systems:

```
size-in-blocks(SYS$SYSTEM:PAGEFILE.SYS)
= size-in-pages(physical-memory) * blocks-per-page
+ number-of-error-log-buffers * blocks-per-buffer
+10
+ value of the system parameter RSRVPAGCNT
```

where:

<i>size-in-pages</i>	Is the size of physical memory, in pages. Use the DCL command SHOW MEMORY to determine the total size of physical memory on your system.
<i>blocks-per-page</i>	Is the number of blocks per page of memory. On Alpha systems, calculate the number of blocks per page of memory by dividing the system's page size by 512 (the size of a block). Use the following commands: <pre>\$ PAGESIZE==F\$GETSYI ("PAGE SIZE") \$ BLOCKSPERPAGE=PAGESIZE/512 \$ SHOW SYMBOL BLOCKSPERPAGE</pre>
<i>number-of-error-log-buffers</i>	Is the value of the system parameter ERRORLOGBUFFERS. This parameter sets the number of error log buffers to permanently allocate in memory.
<i>blocks-per-buffer</i>	Is the value of the system parameter ERLBUFFERPAGES. This parameter sets the number of pagelets (blocks) of memory in each buffer.
RSRVPAGCNT	Is the value of the RSRVPAGCNT special system parameter.

Caution

This formula calculates only the minimum size requirement for saving a dump in the system's primary page file. For most systems, the page file must be larger than this to avoid hanging the system. For more information about calculating the page file size, see Section 16.4.3.

16.4.2 Calculating Error Log Dump File Size

These calculations differ on OpenVMS VAX and Alpha systems:

16.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

On Alpha Systems

On Alpha systems, the AUTOGEN command procedure calculates the appropriate size of your error log dump file. However, to calculate the size of the file manually, use the following formula, which calculates the file size required to hold all the error log buffers:

```
size-in-blocks(SYSS$SYSTEM:SYS$ERRLOG.DMP)
= number-of-error-log-buffers * blocks-per-buffer
+ 2
```

where:

<i>number-of-error-log-buffers</i>	Is the value of the system parameter ERRORLOGBUFFERS. This parameter sets the number of error log buffers that are permanently allocated in memory.
<i>blocks-per-buffer</i>	Is the value of the system parameter ERLBUFFERPAGES. This parameter sets the number of pagelets (blocks) of memory in each buffer.

On VAX Systems

On VAX systems, the size of the error log dump file depends on your DOSD:

- Without DOSD, error logs are written to the dump file SYSDUMP.DMP on the system disk. See Section 16.4.1 for sizing information.
- With DOSD:
 - Both error logs and system memory are written to the SYSDUMP.DMP file on the DOSD disk. See Section 16.4.1 for sizing information.
 - In addition, error logs are written to the stub system dump file SYSDUMP.DMP on the system disk. Because this is a fixed-length file of 2048 blocks, no sizing calculations are required.

16.4.3 Calculating Page File Size

Sufficient page file space is critical to system performance. The AUTOGEN command procedure calculates an appropriate size for your page file space. The size calculated by AUTOGEN should be sufficient. However, to manually calculate the size for page file space, use one of the following formulas.

On VAX Systems

On VAX systems, use the following formula:

```
size-in-blocks (total for all page files on the system)
= size-of-average-process (in pages)
* maximum-number-of-processes
```

- The *size-of-average-process* is the value of the average virtual size of the process. Use the following command to find it:


```
$ SHOW PROCESS/CONTINUOUS/ID=pid
```

Specify this value in pages.
- The *maximum-number-of-processes* is the value of the MAXPROCESSCNT system parameter.

If the result of the formula is less than VIRTUALPAGECNT, use the value of VIRTUALPAGECNT instead.

Managing Page, Swap, and Dump Files

16.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

To determine a system's virtual page count, enter the following command:

```
$ WRITE SYS$OUTPUT F$GETSYI ("VIRTUALPAGECNT")
```

On Alpha Systems

On Alpha systems, no simple formula can be given because of the wide variation possible in memory size and usage. You can use the following formula for systems up to 512MB:

```
size-in-blocks (total for all page files on the system)
= size-of-average-process (in pages)
* blocks-per-page
* maximum-number-of-processes
```

- The *size-of-average-process* is the value of the average virtual size of the process. Use the following command to find it:

```
$ SHOW PROCESS/CONTINUOUS/ID=pid
```

Specify this value in pages.

- Calculate *blocks-per-page* by dividing the system page size by 512 (pagelet size). For example, a system with a page size of 8192 has 16 pagelets per page.

To determine a system's page size, enter the following command:

```
$ WRITE SYS$OUTPUT F$GETSYI ("PAGE_SIZE")
```

- The *maximum-number-of-processes* is the value of the MAXPROCESSCNT system parameter.

For systems over 512MB, follow the steps described below for monitoring page file usage (Section 16.4.3.2), and make adjustments as necessary.

16.4.3.1 Representing Page File Size

The page file size you calculate can be represented in one of the following ways:

- In the primary page file only
- Distributed across primary and secondary page files
- If you have removed the primary page file in SYS\$SYSTEM, distributed across a number of secondary page files

16.4.3.2 Monitoring Page File Usage

Once you determine an initial size for your page file or files (either with AUTOGEN or manually), monitor page file usage by executing AUTOGEN with the following command:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS TESTFILES FEEDBACK
```

With this command, AUTOGEN writes page file usage and size recommendations to the feedback report AGEN\$PARAMS.REPORT. (For more information about AUTOGEN and the feedback report, see Section 15.4 and Section 15.4.2.) The DCL command SHOW MEMORY/FILES also displays file usage, as explained in Section 16.3.

Keep page file usage less than half the size of the page file or files. If a paging file starts to fill to the point where system performance is being affected, a message is printed on the console terminal. If this happens, increase the size of your page file or files or install additional files.

16.4 Manually Calculating Appropriate Sizes for Dump, Page, and Swap Files

Note

Your system resources and work load affect the required size of your page file. You should be familiar with your system resources and work load. For more information, refer to the *OpenVMS Performance Management*.

16.4.3.3 Limiting Page File Space

Limit the amount of page file space consumed by user programs by using the /PGFLQUOTA qualifier of the AUTHORIZE commands ADD and MODIFY. (Refer to the AUTHORIZE section in the *OpenVMS System Management Utilities Reference Manual* for more information.) Do not reduce the value of /PGFLQUOTA below 1024. Size requirements of the page file vary widely, depending on user applications.

16.4.4 Calculating Swap File Size

Sufficient swap file space is critical to system performance. The AUTOGEN command procedure calculates an appropriate size for your swap file space. To manually calculate the size for swap file space, use the following formula:

```
size-in-blocks (total for all swap files on the system)
= maximum-number-of-processes
* average-working-set-quota-of-processes-on-system
```

where:

<i>maximum-number-of-processes</i>	Is the value of the MAXPROCESSCNT system parameter.
<i>average-working-set-quota-of-processes-on-system</i>	Is the average value of the WSQUOTA limit for processes running on the system. On VAX systems, specify the value in pages. On Alpha systems, specify the value in pagelets.

16.4.4.1 Representing Swap File Size

The size you calculate can be represented in any of the following ways:

- In the primary swap file only
- Distributed across primary and secondary swap files
- If you have removed the primary swap file in SYS\$SYSTEM, distributed across a number of secondary swap files

16.4.4.2 Monitoring Swap File Usage

Once you have determined an appropriate size for swap file space (either manually or with AUTOGEN), monitor swap file usage with the DCL command SHOW MEMORY/FILES as explained in Section 16.3. Keep at least one-third of the swap file space unused; otherwise, system performance can be severely affected.

Note

Your system resources and work load affect the required size of your swap file. You should be familiar with your system resources and work load. For more information, refer to the *OpenVMS Performance Management*.

16.5 Minimizing System Dump File Size When Disk Space Is Insufficient

In certain system configurations, it might be impossible to preserve the entire contents of memory in a disk file. For instance, a large memory system might not be able to supply enough disk space for a full memory dump. If your system attempts to save all of memory but the dump file is too small to accommodate the entire dump, the System Dump Analyzer utility (SDA) might not be able to analyze the dump.

On VAX systems, insufficient dump space would also prevent the Crash Log Utility Extractor (CLUE) from being able to analyze the dump.

Options for Minimizing System Dump File Size

Use one of the following options to minimize the size of the system dump file when disk space is insufficient:

- Selective dumps

To preserve those portions of memory that contain information most useful in determining the causes of system failures, you can use selective system dumps. Table 16–1 defines physical and selective dumps.

Table 16–3 compares the information available in physical and selective system dump files.

Table 16–3 Comparison of Physical and Selective System Dump Files

Type	Available Information	Unavailable Information
Physical dump (also known as full dump)	Complete contents of physical memory in use, stored in order of increasing physical address and error log buffers.	Contents of paged-out memory at the time of the crash.
Selective dump	System page table, global page table, system space memory, error log buffers, and process and control regions (plus global pages) for all saved processes.	Contents of paged-out memory at the time of the crash, process and control regions of unsaved processes, and memory not mapped by a page table.

To direct your system to save selective system dumps, set bit 0 of the system parameter DUMPSTYLE. System parameters and their values are in an appendix of the *OpenVMS System Management Utilities Reference Manual*. For information about how to change system parameter values, see Section 15.5.

- Compressed dumps

On Alpha systems, if you set bit 3 of the DUMPSTYLE system parameter, OpenVMS writes the physical or selective system dump in a compressed format. (The exact amount of the compression depends on system use, but the typical compressed dump is about 60 percent of its original size.) If you use compressed dumps, AUTOGEN sizes the system dump file at 2/3 of its uncompressed size.

- Dump off system disk (DOSD)

If you set bit 2 of the DUMPSTYLE system parameter and meet all the other requirements, OpenVMS writes the system dump to a disk other than the system disk. For details, see Section 16.7.

16.5 Minimizing System Dump File Size When Disk Space Is Insufficient

Section 16.5.1 discusses the order in which information is written to a selective system dump on Alpha and VAX systems. Section 16.5.2 discusses how this order can be fine-tuned on Alpha systems.

16.5.1 Understanding the Order of Information in a Selective System Dump

The following lists show the order in which information is written to selective dumps on VAX and Alpha systems.

On VAX systems, information is written to selective dumps in the following order:

1. System page table (SPT)
2. System space (including process page tables, page frame number (PFN) database, and global page table (GPT))
3. Global pages that appear in the working set of any process
4. Processes resident at the time of the crash:
 - a. Current process on crash CPU
 - b. Predefined processes (hardcoded into BUGCHECK)
 - c. Current processes on other CPUs
 - d. Other processes resident at the time of the crash, in order by process index

On Alpha systems, information is written to selective system dumps in the following order:

1. Page table (PT) space for shared addresses (S0/S1/S2)
2. S0/S1 space
3. S2 space
4. Any system space pages (P1, S0/S1, S2) that have been replicated for performance reasons, where the contents of the replicated page is different from the original
5. Memory map pages for Galaxy shared memory regions, if appropriate
6. Key processes:
 - a. Current process on crash CPU
 - b. Swapper
 - c. Current processes on CPUs that have failed to record their crash state
 - d. Current processes on other CPUs
 - e. Site-specific priority processes (see next section)
 - f. Compaq-defined priority processes (hardcoded into BUGCHECK):
 - MSCPmount
 - AUDIT_SERVER
 - NETACP
 - NET\$ACP
 - REMACP
 - LES\$ACP
7. Any processes in a resource or miscellaneous wait state (for example, RWAST)
8. Key global pages (those that appear in the working set of any key process)

Managing Page, Swap, and Dump Files

16.5 Minimizing System Dump File Size When Disk Space Is Insufficient

9. Other processes (the non-key processes) resident at the time of the crash, in order by process index
10. Remaining global pages that appear in the working set of any non-key process

Note that on Alpha systems, processes are dumped in two stages: the page tables for the process first, and then the body of the process.

Usage Notes on VAX and Alpha Systems

On both VAX and Alpha platforms, no process is dumped twice. For example, on Alpha systems, if the current process is the Swapper, it is dumped only once.

Similarly, on Alpha systems, no global page is dumped twice. Therefore, if a page in the working set of a key process is dumped in the "Key global pages" section, it is not dumped again later just because it is also in the working set of a non-key process.

16.5.2 Fine-Tuning the Order That Processes Are Written in Selective System Dumps (Alpha Only)

On Alpha systems, a set of processes, known as **key processes**, are dumped immediately following PT, S0/S1, and S2, including transition pages that link back to the process. The system manager can designate additional processes to be treated as key processes. These processes have priority over other processes in a dump, thus ensuring that the selected processes are successfully written when the dump file is too small to contain all processes.

How to Perform This Task

To designate the order of processes in a dump, follow these steps:

1. Copy the file `SYS$SYSTEM:SYS$DUMP_PRIORITY.TEMPLATE` to `SYS$SYSTEM:SYS$DUMP_PRIORITY.DAT`.
2. Following the instructions in the file, edit the `.DAT` file to contain a list of the processes to be dumped early in the dump.
3. Run the image `SYS$SYSTEM:SYS$SET_DUMP_PRIORITY.EXE`. Note that the image is automatically run during system startup if the data file `SYS$SYSTEM:SYS$DUMP_PRIORITY.DAT` exists.

You can edit the data file and run the image at any time the system is running. Therefore, if a process is hung, the system manager can designate the process as a priority process and then force a crash.

16.6 Writing the System Dump File to the System Disk

If you have more than one path to the system disk, or the system disk is a shadow set with multiple members, you must take additional steps to ensure that a system dump can be written to the system disk.

16.6.1 System Dump to System Disk on Alpha

If there is more than one path to the system disk, the console environment variable `DUMP_DEV` must describe all paths to the system disk. This ensures that if the original boot path becomes unavailable because of failover, the system can still locate the system disk and write the system dump to it.

If the system disk shadow set has multiple members, the console environment variable `DUMP_DEV` must describe all paths to all members of the shadow set. This ensures that if the master member changes, the system can still locate the master member and write the system dump to it.

Managing Page, Swap, and Dump Files

16.6 Writing the System Dump File to the System Disk

If you do not define `DUMP_DEV`, the system can write a system dump only to the physical disk used at boot time using only the same physical path used at boot time. For instructions on setting `DUMP_DEV`, see Section 16.7.1.

Certain configurations (for example, those using Fibre Channel disks) may contain more combinations of paths to the system disk than can be listed in `DUMP_DEV`. In that case, Compaq recommends that you include in `DUMP_DEV` all paths to what is usually the master member of the shadow set, because shadow set membership changes occur less often than path changes.

You can write the system dump to an alternate disk (see Section 16.7.1), but when doing so you must still define a path to the system disk for writing error logs. Therefore, `DUMP_DEV` should contain all paths to the system disk in addition to the paths to the alternate dump disk.

If there are more paths than `DUMP_DEV` can contain, Compaq recommends that you define all paths to the dump disk and as many paths as possible (but at least one) to the system disk. Note that the system disk paths must come last in the list.

16.6.2 System Dump to System Disk on VAX

To ensure that the system can locate the system disk and write the system dump to it when there is more than one path to the system disk, or when the system disk shadow set has multiple members, you must follow the platform-specific instructions regarding booting. On some VAX systems, you must set appropriate register values; on other VAX systems, you must set specific environment variables. See the upgrade and installation supplement for your VAX system for details.

Note that if the system has multiple CI star couplers, the shadow set members must all be connected through the same star coupler.

16.7 Writing the System Dump File to an Alternate Disk

You can write the system dump file to a device other than the system disk (DOSD) on OpenVMS systems. This is especially useful in large-memory systems and in clusters with common system disks where sufficient disk space is not always available on one disk to support customer dump file requirements.

Requirements for DOSD are somewhat different on VAX and Alpha systems. On both systems, however, you must correctly enable the `DUMPSTYLE` system parameter to enable the bugcheck code to write the system dump file to an alternate device.

The following sections describe the requirements for DOSD on Alpha and VAX systems.

16.7.1 DOSD Requirements on Alpha Systems

On Alpha systems, DOSD has the following requirements:

- The dump device directory structure must resemble the current system disk structure. The `[SYSn.SYSEXE]SYSDUMP.DMP` file will reside there, with the same boot time system root.

Use `AUTOGEN` to create this file. In the `MODPARAMS.DAT` file, the following symbol prompts `AUTOGEN` to create the file:

```
DUMPFILE_DEVICE = "$nnn$ddcuuuu"
```

You can enter a list of devices.

Managing Page, Swap, and Dump Files

16.7 Writing the System Dump File to an Alternate Disk

- The dump disk must have an ODS-2 or ODS-5 file structure.
- The dump device cannot be part of a volume set.
- Although not a requirement, Compaq recommends that you mount the dump device during system startup. If the dump device is mounted, it can be accessed by CLUE and AUTOGEN and for the analysis of crash dumps. For best results, include the MOUNT command in SYS\$MANAGER:SYCONFIG.COM.
- For the Crash Log Utility Extractor (CLUE) to support DOSD, you must define the logical name CLUE\$DOSD_DEVICE to point to the dump file to be analyzed after a system crash. For instructions, refer to Section 16.9.
- The dump device cannot be part of a shadow set unless it is also the system device and the master member of the shadow set.
- Use the following format to specify the dump device environment variable DUMP_DEV at the console prompt:

```
>>> SET DUMP_DEV device-name[...]
```

Note

On DEC 3000 series systems, the following restrictions on the use of the DUMP_DEV environment variable exist:

- This variable is not preserved across system power failures because DEC 3000 series systems do not have enough nonvolatile RAM to save the contents of the file. You must reset the DUMP_DEV variable after a power failure. (DUMP_DEV is preserved across all other types of restarts and bootstraps, however.)
- You cannot clear DUMP_DEV (except by power-cycling the system).
- You must use console firmware Version 6.0 or greater because earlier versions do not provide support for DUMP_DEV.

On some CPU types, you can enter only one device; on other CPU types, you can enter a list of devices. The list can include various alternate paths to the system disk and the dump disk.

By specifying an alternate path with DUMP_DEV, the disk can fail over to the alternate path when the system is running. If the system crashes subsequently, the bugcheck code can use the alternate path by referring to the contents of DUMP_DEV.

When you enter a list of devices, however, paths to the system disk must come last.

How to Perform This Task

To designate the dump device with the DUMP_DEV environment variable, and enable the DUMPSTYLE system parameter, follow these steps:

1. Display the value of BOOTDEF_DEV; for example:

```
>>> SHOW BOOTDEF_DEV
BOOTDEF_DEV          dub204.7.0.4.3,dua204.4.0.2.3
```

Managing Page, Swap, and Dump Files

16.7 Writing the System Dump File to an Alternate Disk

2. Display the devices on the system as follows:

```
>>> SHOW DEVICES

Resetting IO subsystem...

dua204.4.0.2.3    $4$DUA204 (RED70A)      RA72
dua206.4.0.2.3    $4$DUA206 (RED70A)      RA72
dua208.4.0.2.3    $4$DUA208 (RED70A)      RA72

polling for units on cixcd1, slot 4, xmi0...

dub204.7.0.4.3    $4$DUA204 (GRN70A)      RA72
dub206.7.0.4.3    $4$DUA206 (GRN70A)      RA72
dub208.7.0.4.3    $4$DUA208 (GRN70A)      RA72
>>>
```

In this example:

- DUA204 is the system disk device.
 - DUA208 is the DOSD device.
3. To provide two paths to the system disk, with the dump disk as DUA208 (also with two paths), set DUMP_DEV as follows:

```
>>> SET DUMP_DEV dua208.4.0.2.3,dub208.7.0.4.3,dub204.7.0.4.3,dua204.4.0.2.3
```

In this example, dua208.4.0.2.3 and dub208.7.0.4.3 are paths to the dump device; dub204.7.0.4.3 and dua204.4.0.2.3 are paths to the boot device.

Note

The system chooses the first valid device that it finds in the list as the dump device. Therefore, the dump disk path entries must appear before the system disk entries in the list.

4. Display all environment variables on the system by entering the SHOW * command; for example:

```
>>> SHOW *

auto_action          HALT
baud_                 9600
boot_dev             dua204.4.0.2.3
boot_file
boot_osflags         0,0
boot_reset           ON
bootdef_dev          dub204.7.0.4.3,dua204.4.0.2.3
booted_dev           dua204.4.0.2.3
booted_file
booted_osflags       0,0
cpu                  0
cpu_enabled          ff
cpu_primary          ff
d_harderr            halt
d_report             summary
d_softerr            continue
dump_dev             dua208.4.0.2.3,dub208.4.0.4.3,dub204.7.0.4.3,dua204.4.0.2.3
enable_audit         ON
interleave           default
language             36
pal                  V5.48-3/01.35-2
prompt               >>>
stored_argc          2
```

Managing Page, Swap, and Dump Files

16.7 Writing the System Dump File to an Alternate Disk

```
stored_argv0      B
stored_argv1      dua204.4.0.2.3
system_variant    0
version           T4.3-4740 Jun 14 2000 15:16:38
>>>
```

5. Enable the DOSD bit of the DUMPSTYLE system parameter by setting bit 2. For example, enter the value of 4 at the SYSBOOT> prompt to designate an uncompressed physical dump to an alternate disk with minimal console output:

```
>>> BOOT
SYSBOOT> SET DUMPSTYLE 4
```

The *OpenVMS System Management Utilities Reference Manual* and online help contain details about the DUMPSTYLE system parameter.

Note

The error log dump file is always created on the system disk so that error log buffers can be restored when the system is rebooted. This file is not affected by setting the DUMPSTYLE system parameter or the DUMP_DEV environmental variable.

16.7.2 DOSD Requirements on VAX Systems

On VAX systems, DOSD has the following requirements:

- The system must be connected directly to, and must boot from, CI controllers.
- The dump device must physically connect to the same two HSx CI controllers as the boot device. These two controllers must be connected through the same CI star coupler.
- The dump device directory structure must resemble the current system disk structure. The [SYS*n*.SYSEXE]SYSDUMP.DMP file will reside there, with the same boot time system root.

Use AUTOGEN to create this file. In the MODPARAMS.DAT file, the following symbol prompts AUTOGEN to create the file:

```
DUMPPFILE_DEVICE = "$nnn$ddcuuuu"
```

You can list only one device.

- The volume label can be up to 12 characters long. The ASCII string DOSD_DUMP must be part of this volume label. For example, valid volume labels are DOSD_DUMP, DOSD_DUMP_12, 12_DOSD_DUMP. The label is read and retained in a memory boot data structure.
- The dump device cannot be part of a volume set. Compaq strongly recommends that the dump device also not be part of a shadow set.
- The dump device cannot be MSCP unit zero (0); only units 1 to 4095 (1—FFF) are supported.

You can designate the dump device as follows:

- On VAX 7000 configurations, by using bits 16 through 27 of the DUMPSTYLE system parameter. Note that the DUMP_DEV environment variable that is provided on VAX 7000 configurations is not used by OpenVMS VAX.

Managing Page, Swap, and Dump Files

16.7 Writing the System Dump File to an Alternate Disk

- On configurations other than the VAX 7000, by using bits 16 through 27 of register 3 (R3). You can use this portion of the register to specify the dump device.

The *OpenVMS System Management Utilities Reference Manual* and online help contain details about the DUMPSTYLE system parameter.

Note

To restore error log buffers when the system is rebooted after a system crash, the error logs must be saved on the system disk. For this purpose, AUTOGEN creates a SYSDUMP.DMP file on the system disk; the file is large enough to contain the maximum size of error log buffers.

16.8 Using SDA to Analyze the Contents of a Crash Dump

The System Dump Analyzer utility (SDA) lets you interpret the contents of the system dump file to investigate the probable causes of the crash. For information about analyzing a crash dump, refer to the *OpenVMS VAX System Dump Analyzer Utility Manual* or the *OpenVMS Alpha System Analysis Tools Manual*.

If your system fails, use SDA to make a copy of the system dump file written at the time of the failure and contact your Compaq support representative. For information about copying the system dump file, see Section 16.12.

16.9 Using SDA CLUE Commands to Analyze Crash Dump Files (Alpha Only)

SDA CLUE (Crash Log Utility Extractor) commands automate the analysis of crash dumps and maintain a history of all fatal bugchecks on a standalone system or cluster. You can use SDA CLUE commands in conjunction with SDA to collect and decode additional dump file information not readily accessible through standard SDA. You can also use SDA CLUE with Dump Off System Disk (DOSD) to analyze a system dump file that resides on a disk other than the system disk.

16.9.1 Understanding CLUE (Alpha Only)

On Alpha systems, SDA is automatically invoked by default when you reboot the system after a system failure. To better facilitate crash dump analysis, SDA CLUE commands automatically capture and archive summary dump file information in a CLUE listing file.

A startup command procedure initiates commands that:

- Invoke SDA
- Issue an SDA CLUE HISTORY command
- Create a listing file called CLUE\$nodename_ddmmyy_hhmm.LIS

The CLUE HISTORY command adds a one-line summary entry to a history file and saves the following output from SDA CLUE commands in the listing file:

- Crash dump summary information
- System configuration
- Stack decoder
- Page and swap files

Managing Page, Swap, and Dump Files

16.9 Using SDA CLUE Commands to Analyze Crash Dump Files (Alpha Only)

- Memory management statistics
- Process DCL recall buffer
- Active XQP processes
- XQP cache header

The contents of this CLUE list file can help you analyze a system failure.

If these files accumulate more space than the threshold allows (default 5000 blocks), the oldest files are deleted until the threshold limit is reached. This can also be customized using the CLUE\$MAX_BLOCK logical name.

To inhibit the running of CLUE at system startup, define the logical CLUE\$INHIBIT in the SYLOGICALS.COM file as /SYS TRUE.

It is important to remember that CLUE\$nodename_ddmmyy_hhmm.LIS contains only an overview of the crash dump and does not always contain enough information to determine the cause of the crash. If you must do an in-depth analysis of the system crash, Compaq recommends that you always use the SDA COPY command to save the dump file.

16.9.2 Displaying Data Using SDA CLUE Commands (Alpha Only)

Invoke CLUE commands at the SDA prompt as follows:

```
SDA> CLUE CONFIG
```

CLUE commands provide summary information of a crash dump captured from a dump file. When debugging a crash dump interactively, you can use SDA CLUE commands to collect and decode some additional information from a dump file, which is not easily accessible through standard SDA. For example, CLUE can quickly provide detailed XQP summaries.

You can also use CLUE commands interactively on a running system to help identify performance problems.

You can use all CLUE commands when analyzing crash dumps; the only CLUE commands that are not allowed when analyzing a running system are CLUE CRASH, CLUE ERRLOG, CLUE HISTORY, and CLUE STACK.

Refer to *OpenVMS Alpha System Analysis Tools Manual* for more information about using SDA CLUE commands.

16.9.3 Using SDA CLUE with Dump Off System Disk (Alpha Only)

Dump off system disk (DOSD) allows you to write the system dump file to a device other than the system disk. For SDA CLUE to be able to correctly find the dump file to be analyzed after a system crash, perform the following steps:

1. Modify the command procedure SYS\$MANAGER:SYCONFIG.COM to add the system logical name CLUE\$DOSD_DEVICE to point to the device where the dump file resides. You need to supply only the physical or logical device name without a file specification.
2. Modify the command procedure SYS\$MANAGER:SYCONFIG.COM to mount systemwide the device where the dump file resides. Otherwise, SDA CLUE cannot access and analyze the dump file.

16.9 Using SDA CLUE Commands to Analyze Crash Dump Files (Alpha Only)

In the following example, the dump file is placed on device `3DUA25`, with the label `DMP$DEV`. You need to add the following commands to `SYS$MANAGER:SYCONFIG.COM`:

```
$mount/system/noassist $3$dua25: dmp$dev dmp$dev
$define/system clue$dosd_device dmp$dev
```

16.10 Using CLUE to Obtain Historical Information About Crash Dumps (VAX Only)

On VAX systems, the Crash Log Utility Extractor (CLUE) displays the contents of a **crash history file**. By examining the contents of the crash history file, you can understand and resolve the issues responsible for failures (crashes), and you might also obtain other useful data.

16.10.1 Understanding CLUE (VAX Only)

The crash history file, which is created and updated by CLUE, contains key parameters from crash dump files. Unlike crash dumps, which are overwritten with each system failure and are therefore typically available only for the most recent failure, the crash history file is a permanent record of system failures.

After a system fails and physical memory is copied to the crash dump file, CLUE automatically appends the relevant parameters to the file `CLUE$OUTPUT:CLUE$HISTORY.DATA` when the system is restarted. The remainder of this section describes how you can use CLUE to display the data it has collected; reference information about CLUE is available in the *OpenVMS System Management Utilities Reference Manual*.

Note

The history file typically grows by about 10 to 15 blocks for each entry. You can limit the number of entries in the binary file by defining the logical name `CLUE$MAX_ENTRIES` to be the maximum number desired. When this number is reached, the oldest entries are deleted from the history file.

By default, operator shutdowns are recorded in the history file. You can exclude information from operator shutdowns in the history file by defining the logical name `CLUE$EXCLUDE_OPERS` as being `TRUE`, for example by including the following line in `SYS$MANAGER:SYSTARTUP_VMS.COM`:

```
$ DEFINE /SYSTEM CLUE$EXCLUDE_OPERS TRUE
```

16.10.2 Displaying Data Using CLUE (VAX Only)

To display data using CLUE, you must first define the following symbol:

```
$ CLUE := $CLUE
```

After defining the symbol, you can use CLUE to display information by entering the following command:

```
$ CLUE/DISPLAY
CLUE_DISPLAY>
```

Managing Page, Swap, and Dump Files

16.10 Using CLUE to Obtain Historical Information About Crash Dumps (VAX Only)

At the CLUE_DISPLAY> prompt, you can issue commands to perform the following actions:

- Use the DIRECTORY command to list failures that have occurred since a specified date, failures of a particular type, failures that contain a specified module, and failures that have a specified offset.

For example, you can list all the failures in the history file using the DIRECTORY command, as follows:

```
CLUE_DISPLAY> DIRECTORY
```

- Use the SHOW command to generate information similar to that obtained from certain commands in the System Dump Analyzer utility (SDA).

For example, if you wanted complete information about the crash listed as crash number 7, the following SHOW command would provide the information:

```
CLUE_DISPLAY> SHOW ALL 7
```

- Use the EXTRACT command to write the data from an entry to a file.

For example, the following command writes the data from entry number 7 in the crash history file to a file named 15MAYCRASH.TXT:

```
CLUE_DISPLAY> EXTRACT 7/OUTPUT=15MAYCRASH.TXT
```

For more information about CLUE commands, refer to the *OpenVMS System Management Utilities Reference Manual*.

16.11 Saving the Contents of the System Dump File After a System Failure

If the system fails, it overwrites the contents of the system crash dump file and the previous contents are lost. For this reason, ensure that your system automatically analyzes and copies the contents of the system dump file each time the system reboots.

On Alpha systems, SDA is invoked by default during startup, and a CLUE list file is created. Generated by a set sequence of commands, the CLUE list file contains only an overview of the crash and might not provide enough information to determine the cause of the crash. Compaq, therefore, recommends that you always copy the system dump file.

Refer to the *OpenVMS Alpha System Analysis Tools Manual* for information about modifying your site-specific command procedure to execute additional commands such as SDA COPY upon startup after a system failure.

On VAX systems, modify the site-specific startup command procedure SYSTARTUP_VMS.COM so that it invokes the System Dump Analyzer utility (SDA) when the system is booted.

Be aware of the following facts:

- When invoked from the site-specific startup procedure in the STARTUP process, SDA executes the specified commands only if the system is booting immediately after a system failure. If the system is rebooting after it was shut down with SHUTDOWN.COM or OPCCRASH.EXE, SDA exits without executing the commands.

16.11 Saving the Contents of the System Dump File After a System Failure

- Although you can use the DCL command COPY to copy the dump file, the SDA command COPY is preferable because it copies only the blocks occupied by the dump and it marks the dump file as copied. The SDA COPY command is preferable also when the dump was written into the paging file, SYS\$SYSTEM:PAGEFILE.SYS, because the SDA COPY command releases to the pager those pages occupied by the dump. For more information, see Section 16.13.
- Because a system dump file can contain privileged information, protect copies of dump files from world read access. For more information about file protection, refer to the *OpenVMS Guide to System Security*.
- System dump files have the NOBACKUP attribute, so the Backup utility (BACKUP) does not copy them unless you use the qualifier /IGNORE=NOBACKUP when invoking BACKUP. When you use the SDA command COPY to copy the system dump file to another file, the operating system does not automatically set the new file to NOBACKUP. If you want to set the NOBACKUP attribute on the copy, use the SET FILE command with the /NOBACKUP qualifier as described in the *OpenVMS DCL Dictionary*.

Example

The SDA command COPY in the following example saves the contents of the file SYS\$SYSTEM:PAGEFILE.SYS and performs some analysis of the file. Note that the COPY command is the final command because the blocks of the page file used by the dump are released as soon as the COPY command completes, and can be used for paging before any other SDA commands can be executed.

```
$ !
$ !      Print dump listing if system just failed
$ !
$ ANALYZE/CRASH_DUMP SYS$SYSTEM:PAGEFILE.SYS
   SET OUTPUT DISK1:SYSDUMP.LIS      ! Create listing file
   READ/EXECUTIVE                    ! Read in symbols for kernel
   SHOW CRASH                        ! Display crash information
   SHOW STACK                        ! Show current stack
   SHOW SUMMARY                      ! List all active processes
   SHOW PROCESS/PCB/PHD/REG          ! Display current process
   COPY SYS$SYSTEM:SAVEDUMP.DMP      ! Save system dump file
   EXIT
$ SET FILE/NOBACKUP SYS$SYSTEM:SAVEDUMP.DMP
```

16.12 Copying System Dump Files to Tape or Disk

If your system fails, make a copy of the contents of the system dump file and contact your Compaq support representative. You can use the Backup utility (BACKUP) to create save sets containing system dump files on magnetic tape or disk. However, when using BACKUP to copy system dump files, you must specify the /IGNORE=(NOBACKUP,INTERLOCK) qualifier for the following reasons:

- By default, the system dump file has the NOBACKUP attribute, so it is not copied unless you specify /IGNORE=NOBACKUP.
- The system keeps an open channel to the system dump file, so the file is not copied unless you specify /IGNORE=INTERLOCK.

For more information about using BACKUP, see Chapter 11. For information about BACKUP commands, see the BACKUP section in the *OpenVMS System Management Utilities Reference Manual*.

Managing Page, Swap, and Dump Files

16.12 Copying System Dump Files to Tape or Disk

Compaq recommends that you use the following procedure to copy your system dump file:

1. Use the SDA command COPY to make a copy of the system dump file.
2. Use BACKUP to save that copy on tape or disk.

This procedure avoids issues with BACKUP qualifiers and reduces the amount of data written to tape or disk because the SDA command COPY copies only blocks in the system dump file that are actually used.

16.13 Freeing Dump Information from the Page File

If you use SYS\$SYSTEM:PAGEFILE.SYS to store a system crash dump, you must later free the space occupied by the system dump for use by the pager. If you do not, your system might hang because of insufficient paging space.

Section 16.1.1 explains when you might use the page file to store a system crash dump.

16.13.1 Freeing Dump Information on VAX and Alpha Systems

This section contains instructions for freeing dump information from the page file on VAX and Alpha systems.

How to Perform This Task on VAX Systems

On VAX systems, perform the following steps:

1. Invoke the System Dump Analyzer utility (SDA), specifying PAGEFILE.SYS as the target:

```
$ ANALYZE/CRASH_DUMP SYS$SYSTEM:PAGEFILE.SYS
```

2. Enter the SDA command COPY in the following format to copy the dump from SYS\$SYSTEM:PAGEFILE.SYS to another file:

```
COPY dump_filespec
```

For example, to copy the system dump file off the system disk to a file called SAVEDUMP.DMP on DISK\$USER5, enter the following command:

```
SDA> COPY DISK$USER5:[DUMPS]SAVEDUMP.DMP
```

3. Enter the EXIT command to exit SDA.
4. Include the SDA commands entered in steps 1 and 2 in the site-specific startup command procedure SYSTARTUP_VMS.COM to free page space each time the system reboots.

Alternatively, to free the pages in the page file that are taken up by the dump without having to copy the dump elsewhere, enter the ANALYZE/CRASH_DUMP/RELEASE command. This command immediately releases the pages to be used for system paging, effectively deleting the dump. Note that this command does *not* allow you to analyze the dump before deleting it.

Example

The following commands, added to the SYSTARTUP_VMS.COM command procedure, copy the contents of the page file to a file named SAVEDUMP.DMP:

```
$ ANALYZE/CRASH_DUMP SYS$SYSTEM:PAGEFILE.SYS
  COPY DISK$USER5:[DUMPS]SAVEDUMP.DMP
  EXIT
$ SET FILE/NOBACKUP SYS$SYSTEM:SAVEDUMP.DMP
```

Managing Page, Swap, and Dump Files

16.13 Freeing Dump Information from the Page File

How to Perform This Task on Alpha Systems

On Alpha systems, as described in the *OpenVMS Alpha System Analysis Tools Manual*, SDA is automatically invoked by default when the system is rebooted after a system failure.

To automatically save the system dump file, perform the following steps:

1. Create a SYS\$MANAGER:SAVEDUMP.COM file; for example:

```
!
! SDA command file, to be executed as part of the system
! bootstrap from within CLUE.  Commands in this file can
! be used to save the dump file after a system bugcheck, and
! to execute any additional SDA command.
!
READ/EXEC           ! Read in the executive images' symbol tables
SHOW STACK         ! Display the stack
COPY SAVEDUMP.DMP  ! Copy and save system dump file
!
```

2. To point to your site-specific file, add a line such as the following one to the file SYS\$MANAGER:SYLOGICALS.COM:

```
$ DEFINE/SYSTEM CLUE$SITE_PROC SYS$MANAGER:SAVEDUMP.COM
```

In this example, the site-specific file is named SAVEDUMP.COM.

If the logical CLUE\$INHIBIT has been defined, and SDA has not been automatically invoked during system startup, the pages in the page file that are taken up by the dump can be released using the ANALYZE/CRASH_DUMP/RELEASE command. This command immediately releases the pages to be used for system paging, effectively deleting the dump. Note that this command does *not* allow you to analyze the dump before deleting it.

For a discussion of logical names used by CLUE, refer to *OpenVMS Alpha System Analysis Tools Manual*.

16.13.2 Usage Notes for Freeing Information on VAX and Alpha Systems

Because a system dump file can contain privileged information, protect copies of dump files from world read access.

To prevent the system from backing up the complete contents of the file, assign the NOBACKUP attribute to the file with the DCL command SET FILE/NOBACKUP.

Although you can also use the DCL command COPY to copy a dump file, Compaq recommends that you use the SDA command COPY because SDA COPY performs the following actions:

- Copies only the blocks that the dump actually occupies.
- Releases for paging the pages that the dump occupies in the system's page file.

16.14 Installing Page and Swap Files

The system automatically installs the primary page and swap files located in SYS\$SYSTEM. However, other page and swap files are not automatically installed. For this reason, if you create secondary page and swap files, you must also install them with the System Generation utility (SYSGEN). Note that SYSGEN INSTALL commands perform a different function than Install utility (INSTALL) commands.

Managing Page, Swap, and Dump Files

16.14 Installing Page and Swap Files

16.14.1 Installing Interactively

1. Invoke SYSGEN by entering the following command:

```
$ RUN SYS$SYSTEM:SYSGEN
```
2. Enter the SYSGEN command `INSTALL` as follows:
For page files, use the following format:

```
INSTALL file-spec/PAGEFILE
```


For example:

```
SYSGEN> INSTALL DUA2:[PAGE_SWAP]PAGEFILE_1.SYS/PAGEFILE
```


For swap files, use the following format:

```
INSTALL file-spec/SWAPFILE
```


For example:

```
SYSGEN> INSTALL DUA2:[PAGE_SWAP]SWAPFILE_1.SYS/SWAPFILE
```
3. To make sure the files are installed each time the system boots, edit `SYS$MANAGER:SYPAGSWPFILES.COM` to add the commands you entered in step 2. For more information, see Section 16.14.2.

Example

The following example installs page and swap files interactively:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> INSTALL DUA2:[PAGE_SWAP]PAGEFILE_1.SYS/PAGEFILE
SYSGEN> INSTALL DUA2:[PAGE_SWAP]SWAPFILE_1.SYS/SWAPFILE
```

16.14.2 Installing in `SYPAGSWPFILES.COM`

Page and swap files other than `SYS$SYSTEM:PAGEFILE.SYS` and `SYS$SYSTEM:SWAPFILE.SYS` must be reinstalled each time the system boots. You can do this by adding the commands to install the files to the startup command procedure `SYS$MANAGER:SYPAGSWPFILES.COM`. The template file `SYS$MANAGER:SYPAGSWPFILES.TEMPLATE` includes comments that help explain how this file is used.

Before performing this task, you must have created the secondary files, as explained in Section 16.16.

For more information about `SYPAGSWPFILES.COM`, see Section 5.2.3.

You can also use `SATELLITE_PAGE.COM` to install page and swap files on an OpenVMS Cluster satellite node's local disk. `SATELLITE_PAGE.COM` is created when you run `CLUSTER_CONFIG.COM`. For more information about installing page and swap files on a satellite node's local disk, refer to the *OpenVMS Cluster Systems* manual.

How to Perform This Task

1. Invoke any editor to edit `SYS$MANAGER:SYPAGSWPFILES.COM`.
2. If necessary, add a `MOUNT` command for each disk that holds a page or swap file. This is necessary because only the system disk is mounted at the time `SYPAGSWPFILES.COM` is invoked.

For example:

```
$ MOUNT/SYSTEM/NOASSIST DUA2: DISK_SYS2
```

Managing Page, Swap, and Dump Files

16.14 Installing Page and Swap Files

For information about the MOUNT command, refer to the *OpenVMS DCL Dictionary*.

The following commands, inserted before the MOUNT command, are also useful to determine if the disk is available before mounting. Note, however, that if the disk is broken and cannot mount, these commands will cause an infinite loop.

```
$ LOOP1:
$ ON WARNING THEN GOTO LOOP1
$ WAIT 0000 00:00:00.50
$ READY = F$GETDVI("device:", "AVL")
$ IF READY .EQS. "FALSE" THEN GOTO LOOP1
```

where *device:* specifies the device name.

3. Add the following command to invoke SYSGEN:

```
$ RUN SYS$SYSTEM:SYSGEN
```

4. Add commands in the following format to SYPAGSWPFILES.COM to install the files each time the system boots.

For page files, use the following format:

```
INSTALL file-spec/PAGEFILE
```

For example:

```
INSTALL DUA2:[SYSTEM]PAGEFILE_1.SYS/PAGEFILE
```

For swap files, use the following format:

```
INSTALL file-spec/SWAPFILE
```

For example:

```
INSTALL DUA2:[SYSTEM]SWAPFILE_1.SYS/SWAPFILE
```

5. Add an EXIT command to exit SYSGEN:

```
EXIT
```

Example

The following example shows commands you might add to SYPAGSWPFILES.COM to install page and swap files named PAGEFILE_1.SYS and SWAPFILE_1.SYS located on the DUA2: device:

```
$ EDIT SYS$MANAGER:SYPAGSWPFILES.COM
[add the following commands to SYPAGSWPFILES.COM:]
.
.
.
$ MOUNT/SYSTEM/NOASSIST DUA2: DISK_SYS2
$ RUN SYS$SYSTEM:SYSGEN
INSTALL DUA2:[SYSTEM]PAGEFILE_1.SYS /PAGEFILE
INSTALL DUA2:[SYSTEM]SWAPFILE_1.SYS /SWAPFILE
EXIT
```

Managing Page, Swap, and Dump Files

16.15 Removing Page, Swap, and Dump Files

16.15 Removing Page, Swap, and Dump Files

Caution

If you remove a system page, swap, or dump file, do not simply delete the file. The disk might become corrupted if you continue to use the system after you delete the files.

How to Perform This Task

1. Use the `RENAME` command to rename the file to be deleted.
2. Shut down and reboot the system.
3. Delete the file.
4. When you delete a file, make sure you remove from `SYPAGSWPFILES.COM` and `MODPARAMS.DAT` any command lines related to the file.

Example

```
$ RENAME DUA2:[SYSTEM]PAGEFILE_1.SYS; DUA2:[SYSTEM]JUNK.SYS;
$ @SYS$SYSTEM:SHUTDOWN.COM
.
.
.
[SHUTDOWN.COM shuts down and reboots the system]
[When the system reboots, log in]
.
.
.
$ DELETE DUA2:[SYSTEM]JUNK.SYS;
```

16.16 Creating and Modifying Page, Swap, and Dump Files

For performance or disk space reasons, you might want to create system page, swap, and dump files on disks other than the system disk. (Error log dump files, however, must remain on the system disk.)

The following sections explain how to perform this task:

Method	For More Information
Using <code>AUTOGEN</code> (recommended method)	Section 16.16.1
Using <code>SWAPFILES.COM</code> (for primary files only)	Section 16.16.2
Using <code>SYSGEN</code>	Section 16.16.3

16.16.1 Using `AUTOGEN` (Recommended Method)

You can direct `AUTOGEN` to create new system page, swap, and dump files by adding symbols to `MODPARAMS.DAT` to specify the name, location, and size of new files to be created and then running `AUTOGEN`. Before performing this task, you should understand `AUTOGEN` and its parameter file `MODPARAMS.DAT`. For more information about when to use `AUTOGEN`, see Section 15.4. See Section 15.4.4 for information about `MODPARAMS.DAT`.

`AUTOGEN` automatically calculates appropriate sizes for system page, swap, and dump files. It also modifies the files to the appropriate sizes and installs them. You can control sizes calculated by `AUTOGEN` by defining symbols in the file `MODPARAMS.DAT`. For more information, see Section 16.16.1.2.

Managing Page, Swap, and Dump Files

16.16 Creating and Modifying Page, Swap, and Dump Files

How to Perform This Task

To change the sizes of system page, swap, and dump files, execute AUTOGEN in two passes as follows:

1. Enter the following command to invoke a first pass of AUTOGEN. AUTOGEN displays its calculations for system file sizes to SYS\$OUTPUT:

```
$ @SYS$UPDATE:AUTOGEN SAVPARAMS TESTFILES
```

2. If the file sizes displayed in step 1 are inadequate, add symbols to MODPARAMS.DAT to control the size of files as explained in Section 16.16.1.2 and return to step 1.
3. When you are satisfied with the file sizes displayed in step 1, execute a second pass of AUTOGEN using the following command to install the modified system files when the system is rebooted:

```
$ @SYS$UPDATE:AUTOGEN GENPARAMS REBOOT
```

4. Add commands to the site-specific startup command procedure SYPAGSWPFILES.COM to make sure the files are installed each time the system boots. For instructions, see Section 16.14.

16.16.1.1 Controlling the Location of System Page, Swap, and Dump Files

Add the following symbols to MODPARAMS.DAT to specify the names and locations of the page and swap files to be created:

Definition	For Page Files	For Swap Files	For Dump Files
File name and location	PAGEFILE n _NAME = " <i>file-spec</i> "	SWAPFILE n _NAME = " <i>file-spec</i> "	DUMPFIL_ DEVICE = " <i>device</i> "

where:

- n specifies the page or swap file. Refer to the primary page and swap files by specifying a value of 1 for n ; refer to subsequent files by specifying increasingly higher integer values for n . For example, to refer to a secondary page or swap file, specify a value of 2 for n .
- *file-spec* specifies the full file specification of the file to be created, and should be within quotation marks ("").
- *device* specifies the disk name to be used, and should be within quotation marks ("").

16.16.1.2 Controlling the Size of System Page, Swap, and Dump Files in MODPARAMS.DAT

You can add information to the AUTOGEN parameter file MODPARAMS.DAT to control the sizes that AUTOGEN calculates for system page, swap, and dump files. If you do not supply system file size information in MODPARAMS.DAT, AUTOGEN performs default size calculations for page, swap, and dump files.

You can define symbols in MODPARAMS.DAT to specify either of the following items:

Managing Page, Swap, and Dump Files

16.16 Creating and Modifying Page, Swap, and Dump Files

Size to Be Specified	For More Information
<i>Total</i> desired size for all page or swap files on a system (not valid for the system dump files)	Table 16-4
Sizes for <i>individual</i> page, swap, or dump files	Table 16-5

Note

You cannot specify sizes for both total and individual files. AUTOGEN issues a warning if conflicting symbol definitions exist in MODPARAMS.DAT.

For page and swap files, AUTOGEN generally manipulates the primary files SYS\$SYSTEM:PAGEFILE.SYS and SYS\$SYSTEM:SWAPFILE.SYS *only* if you have no other page and swap files. If you have secondary files, AUTOGEN manipulates the secondary files and excludes primary files. However, in some instances, AUTOGEN might modify the size of the primary page and swap files.

On VAX systems, for system dump files, AUTOGEN manipulates the size of only one file: the system dump file on the system disk if no DUMPFIL_E_DEVIC_E is given, or the system dump file on the specified device if DUMPFIL_E_DEVIC_E is specified.

On VAX systems, AUTOGEN always creates a minimal SYSDUMP.DMP file on the system disk for error log buffers if DUMPFIL_E_DEVIC_E is specified.

On Alpha systems, AUTOGEN only manipulates the size of the error log dump file on the system disk.

If you do not want AUTOGEN to change the sizes of the primary files, specify the following symbols in MODPARAMS.DAT:

```
PAGEFILE = 0
SWAPFILE = 0
DUMPFILE = 0
ERRLOGDUMP = 0 ! Alpha only
```

These symbols direct AUTOGEN to ignore the primary page, swap, and dump files when calculating sizes.

If the creation or extension of a system page, swap, or dump file would cause the target disk to become more than 95 percent full, AUTOGEN issues a warning and does not perform the operation.

On Alpha systems, however, the 95 percent rule does not apply to the error log dump file, SYS\$ERRLOG.DMP. This file is created if the disk can hold it.

You can use AUTOGEN to create a page, swap, or dump file that is smaller than the current version of the file. After you have booted and begun using the new file, remember to use the DCL command PURGE to reclaim the disk space from the old version of the file.

To determine the current sizes of installed page and swap files, enter the DCL command SHOW MEMORY/FILES. If you increased the size of any of these files and have not rebooted, this command displays the original sizes. Use the DIRECTORY command to determine the size of dump files.

Managing Page, Swap, and Dump Files

16.16 Creating and Modifying Page, Swap, and Dump Files

Note

AUTOGEN does not change file sizes if you specify a value of 0 or a value that is within 10 percent of the current size.

Table 16–4 lists the symbols you can define in MODPARAMS.DAT to control the *total* size of page file, swap file, system dump file, or error log dump file space.

Table 16–4 Symbols for Controlling the Total Size of Page, Swap, System Dump, or Error Log Dump File Space

Operation	Page File Symbol	Swap File Symbol	Dump File Symbol	Error Log File Symbol
To define the total amount of space	PAGEFILE = n^1	SWAPFILE = n^1	DUMPFIL = n^1	ERRLOGDUMP = n^1
To increase total size	ADD_PAGEFILE = n	ADD_SWAPFILE = n	ADD_DUMPFIL = n	ADD_ERRLOGDUMP = n
To specify maximum total size	MAX_PAGEFILE = n	MAX_SWAPFILE = n	MAX_DUMPFIL = n	MAX_ERRLOGDUMP = n
To specify minimum total size	MIN_PAGEFILE = n	MIN_SWAPFILE = n	MIN_DUMPFIL = n	MIN_ERRLOGDUMP = n

¹ n is the total size, in blocks. If n is 0, the corresponding AUTOGEN section is skipped. For page and swap files, if n is not 0 and no secondary files exist, AUTOGEN applies the value to primary files. If n is not 0, and secondary files exist, AUTOGEN applies any change evenly across all secondary page or swap files but, in most cases, does not change primary files. For dump files, if n is not 0, AUTOGEN applies the value to the dump file on the system disk if no DUMPFIL_DEVICE is given or the dump file on the specified device if a DUMPFIL_DEVICE is given.

Table 16–5 lists the symbols you can define in MODPARAMS.DAT to control the size of *individual* files.

Table 16–5 Symbols for Controlling the Size of Individual Page and Swap Files

Operation	Page File Symbol ¹	Swap File Symbol ¹
To specify file size	PAGEFILE n_SIZE = <i>block-size</i>	SWAPFILE n_SIZE = <i>block-size</i>
To increase file size	ADD_PAGEFILE n_SIZE = <i>block-size</i>	ADD_SWAPFILE n_SIZE = <i>block-size</i>
To specify maximum file size	MAX_PAGEFILE n_SIZE = <i>block-size</i>	MAX_SWAPFILE n_SIZE = <i>block-size</i>
To specify minimum file size	MIN_PAGEFILE n_SIZE = <i>block-size</i>	MIN_SWAPFILE n_SIZE = <i>block-size</i>

¹For n , specify an integer that indicates the page or swap file. Refer to the primary page and swap files by specifying a value of 1 for n ; refer to subsequent files by specifying increasingly higher integer values for n . For example, to refer to a secondary page or swap file, specify a value of 2 for n . For *block-size*, specify the size in blocks.

Examples

- The following line in MODPARAMS.DAT specifies that all page file space should total 100,000 blocks:

```
PAGEFILE = 100000
```

Managing Page, Swap, and Dump Files

16.16 Creating and Modifying Page, Swap, and Dump Files

If you had only a primary page file, the resulting size of that file would be 100,000 blocks. If you had multiple page files, the difference between the total current size and the total new size would be spread across secondary files. For example, if you specified `PAGEFILE = 100000`, the changed page file sizes would be as follows:

File	Original Size (in Blocks)	Resulting Size (in Blocks)
Primary page file	10,000	10,000
Secondary page file 1	30,000	45,000
Secondary page file 2	30,000	45,000

2. To direct AUTOGEN to set the primary page file size to 10,000 blocks, use the symbol definition:

```
PAGEFILE1_SIZE = 10000
```

3. To direct AUTOGEN to create a new secondary swap file named `PAGED$:[PAGESWAP]SWAPFILE.SYS` that holds 30,000 blocks, use the symbol definitions:

```
SWAPFILE2_NAME = "PAGED$:[PAGESWAP]SWAPFILE.SYS"  
MIN_SWAPFILE2_SIZE = 30000
```

16.16.2 Using SWAPFILES.COM

Compaq recommends that you use AUTOGEN to change sizes of system page, swap, and dump files. However, you can use the command procedure `SYS$UPDATE:SWAPFILES.COM` to change the size of *primary* system page, swap, and dump files. `SWAPFILES.COM` shows you the current size of the system page, swap, and dump files before you change the sizes.

If you change the sizes of system page, swap, or dump files, you must edit `MODPARAMS.DAT` to specify the new sizes, as explained in Section 16.16.1.2. If you do not specify the new sizes in `MODPARAMS.DAT`, AUTOGEN resizes the files next time it runs.

The procedure displays the sizes of the current system page, swap, and dump files in `SYS$SYSTEM`, and the amount of space remaining on the system disk. It then allows you to enter new sizes, or keep the existing sizes for these files. If you specify a size that is larger than that of an existing file, the procedure automatically extends the size of a page or dump file. If you specify a smaller size for a system page, swap, or dump file, a new version of the file is created.

How to Perform This Task

1. Enter the following command to invoke the command procedure:

```
$ @SYS$UPDATE:SWAPFILES.COM
```

The system displays the current files found in `SYS$SYSTEM` and their sizes. For example:

```
Current file sizes are:
```

```
Directory SYS$SYSROOT:[SYSEXE]
```

```
PAGEFILE.SYS;1    16384  
SYSDUMP.DMP;1    4128  
SWAPFILE.SYS;1   3072
```

```
Total of 3 files, 23584 blocks.
```

Managing Page, Swap, and Dump Files

16.16 Creating and Modifying Page, Swap, and Dump Files

There are 128741 available blocks on SYS\$SYSDEVICE.

2. In response to the following prompt, type the desired size, in blocks, for the page file. To keep the same size, press Return:
Enter new size for page file:
3. In response to the following prompt, type the desired size, in blocks, for the dump file. To keep the same size, press Return:
Enter new size for system dump file:
4. In response to the following prompt, type the desired size, in blocks, for the swap file. To keep the same size, press Return:
Enter new size for swap file:
5. Shut down and reboot the system to use the new files.
6. After the system reboots, purge obsolete copies of the files. Do not delete the old files until the system reboots.
7. Edit MODPARAMS.DAT to include the new file sizes, as explained in Section 16.16.1.2. If you do not specify the new sizes in MODPARAMS.DAT, AUTOGEN will automatically resize the files the next time it runs.

Example

```
$ @SYS$UPDATE:SWAPFILES.COM
```

To leave a file size at its current value type a carriage return in response to its size prompt. Current file sizes are:

```
Directory SYS$SYSROOT:[SYSEXE]
```

```
PAGEFILE.SYS;1    100000  
SYSDUMP.DMP;1    28000  
SWAPFILE.SYS;1   33000
```

Total of 3 files, 161000 blocks.

There are 128741 available blocks on SYS\$SYSDEVICE.

```
Enter new size for page file: 
```

```
Enter new size for system dump file: 30000
```

```
%SYSGEN-I-EXTENDED, SYS$SYSROOT:[SYSEXE]SYSDUMP.DMP;1 extended
```

```
Enter new size for swap file: 
```

```
*****  
* Please reboot in order for the new files to be used by the system. *  
* After rebooting, purge obsolete copies of the files. *  
* DO NOT delete the old files until after the reboot. *  
*****
```

16.16.3 Using SYSGEN

Compaq recommends that you use AUTOGEN to create and change page, swap, and dump files. AUTOGEN invokes the System Generation utility (SYSGEN) to create or change the files. However, in an emergency, you can use SYSGEN to directly change the size of page, swap and dump files. For example, if you see that page file space is becoming dangerously low, you might use SYSGEN to quickly add page file space to prevent the system from hanging.

Managing Page, Swap, and Dump Files

16.16 Creating and Modifying Page, Swap, and Dump Files

Note

VAX: System parameters SWPFILCNT and PAGFILCNT limit the number of swap and page files that the system installs. Refer to the *OpenVMS System Management Utilities Reference Manual* for more information.

Alpha: OpenVMS Alpha supports a maximum of 254 page and or swap files on a system. System parameters SWPFILCNT and PAGFILCNT are not used on Alpha systems.

How to Perform This Task

1. Determine the location and appropriate size of the files. For information, see Section 16.4.
2. Invoke SYSGEN and enter the CREATE command in the following format:

```
CREATE file-spec/SIZE=block-count
```

where:

file-spec specifies the full file specification.

block-count specifies the size of the file in blocks.

If the file you specify already exists and the size you specify is larger than the existing file, the command extends the existing file. If the file you specify already exists and the size you specify is smaller than the existing file, the command creates a new file of the specified size.

For example, the following command extends the existing, smaller primary page file PAGEFILE.SYS:

```
SYSGEN> CREATE PAGEFILE.SYS/SIZE=100000
```

For more information about the SYSGEN command CREATE, refer to the SYSGEN section in the *OpenVMS System Management Utilities Reference Manual*.

Note

Frequent file creation and deletion can cause the free space on a disk to become severely fragmented. SYSGEN issues a HEADERFULL warning message if it determines that the creation or extension of a system file would cause that file to become fragmented enough to render the system unbootable. If this occurs, Compaq recommends that you back up and restore your system disk to consolidate the free space on the volume into one contiguous area. For more information, see Section 11.17.

After you restore the disk, retry the SYSGEN operation. When SYSGEN issues a warning message, the file might be somewhat larger, but not as large as the value specified in the CREATE command.

3. Use the following table to determine if you should reboot to use the new or modified file:

Managing Page, Swap, and Dump Files

16.16 Creating and Modifying Page, Swap, and Dump Files

Type	Change	Reboot Required?
Primary page, swap, system dump, or error log dump file ¹	New file	Yes
	Extended file	Yes
Secondary page or swap file	New file	No ²
	Extended file	Yes
Alternate (DOSD) dump file (Alpha)	New file	No
	Extended file	No
Alternate (DOSD) dump file (VAX)	New file	Yes
	Extended file	Yes

¹Primary page, swap, and dump files are SYS\$SPECIFIC:[SYSEXE] PAGEFILE.SYS, SWAPFILE.SYS, SYSDUMP.DMP; and SYS\$ERRLOG.DMP.

²Although rebooting the system is unnecessary, you must install secondary files before the system can use them. For more information, see Section 16.14.

4. If you create a new version of the file, purge the old version *after* the system reboots.
5. Add commands to the site-specific startup command procedure SYPAGSWPFILES.COM to make sure the files are installed each time the system boots. For instructions, see Section 16.14.
6. If you do not want AUTOGEN to resize the files according to its calculations, edit MODPARAMS.DAT to specify the sizes of these files. Follow the instructions in Section 16.16.1.2.

Example

The commands in the following example extend the existing files PAGEFILE.SYS, SWAPFILE.SYS, and SYSDUMP.DMP to the specified sizes:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> CREATE PAGEFILE.SYS/SIZE=100000
%SYSGEN-I-EXTENDED, SYS$SYSROOT:[SYSEXE]PAGEFILE.SYS;1 extended
SYSGEN> CREATE SWAPFILE.SYS/SIZE=30000
%SYSGEN-I-EXTENDED, SYS$SYSROOT:[SYSEXE]SWAPFILE.SYS;1 extended
SYSGEN> CREATE SYSDUMP.DMP/SIZE=33000
%SYSGEN-I-EXTENDED, SYS$SYSROOT:[SYSEXE]SYSDUMP.DMP;1 extended
SYSGEN> EXIT
```

16.17 Understanding Process Dumps

When a single process fails but the operating system is still running, the system can create a **process dump** that contains information about the process to assist in determining what caused the process to fail.

By default, process dumps are written to the current default directory of the user. You can override this by defining the logical name SYS\$PROCDMP to identify an alternate directory path. Note that the name of the process dump file is always the same as the name of the main image active at the time the process dump is written, with the file extension .DMP.

On Alpha systems, a process dump is either complete or partial. A **complete process dump** contains all of process space and all process-pertinent data from system space. A **partial process dump** contains only user-readable data from process space and only those data structures from system space that are not

Managing Page, Swap, and Dump Files

16.17 Understanding Process Dumps

deemed sensitive. Privileged or protected data, such as an encryption key in third-party software, might be considered sensitive.

On Alpha systems, you can force a dump to be written for another process with the DCL command `SET PROCESS/DUMP=NOW process-spec`. This command causes the contents of the address space occupied by *process-spec* to be written immediately to the file named *image-name.DMP* in the current directory of *process-spec*.

For more information about the DCL command `SET PROCESS/DUMP`, refer to the *OpenVMS DCL Dictionary: N–Z*.

16.17.1 Understanding Privileged Users and Access to Process Dumps (Alpha Only)

For this discussion, a privileged user is one who satisfies one of the following conditions:

- Has one or more of the privileges `CMKRNL`, `CMEXEC`, `SYSPRV`, `READALL`, or `BYPASS`
- Is a member of a system UIC group (by default `[10,n]` or lower). Such users are treated as though they hold `SYSPRV` privilege.

Holders of `CMKRNL` or `CMEXEC` can write complete process dumps. Holders of any of the other privileges mentioned above can read a process dump wherever it has been written.

In general, nonprivileged users should not be able to read complete process dumps, and by default they cannot do so. However, certain situations require that a nonprivileged user be able to read a complete process dump. Other situations require that a nonprivileged user be able to create a complete process dump but be able to read only a partial process dump.

Rights identifier `IMGDMP$READALL` enables a nonprivileged user to read a complete process dump. Rights identifier `IMGDMP$PROTECT` protects a complete process dump from being read by the nonprivileged user that created the process dump. These rights identifiers are created during the installation of OpenVMS by the image `SYS$SYSTEM:IMGDMP_RIGHTS.EXE`, which is also run automatically during system startup to ensure that these rights identifiers exist with the correct values and attributes.

If these rights identifiers have been deleted, you can run `SYS$SYSTEM:IMGDMP_RIGHTS.EXE` to recreate them. For example:

```
$ RUN SYS$SYSTEM:IMGDMP_RIGHTS
%PROC_DUMP-I-CREATED, rights identifier IMGDMP$READALL successfully created
%PROC_DUMP-I-CREATED, rights identifier IMGDMP$PROTECT successfully created
```

Note that `IMGDMP$READALL` has no attributes, but `IMGDMP$PROTECT` is created with the `RESOURCE` attribute.

16.17.2 Granting Access to Process Dumps (Alpha Only)

To allow a nonprivileged user to write and read complete process dumps, grant the rights identifier `IMGDMP$READALL` to the user. If the `IMGDMP$READALL` rights identifier does not exist, run the image `SYS$SYSTEM:IMGDMP_RIGHTS.EXE` to create it (see Section 16.17.1). Then use `AUTHORIZE` to grant the rights identifier to the user. For example:

Managing Page, Swap, and Dump Files

16.17 Understanding Process Dumps

```
$ DEFINE /USER SYSUAF SYS$SYSTEM:SYSUAF.DAT !if necessary
$ RUN SYS$SYSTEM:AUTHORIZE
UAF> GRANT /IDENTIFIER IMGDMPS$READALL <user>
UAF> EXIT
```

Note that the user must log out and log in again to be able to exercise the rights identifier. A nonprivileged user with rights identifier IMGDMPS\$READALL can read and write complete process dumps without restriction.

16.17.3 Restricting Access to Process Dumps (Alpha Only)

You can allow a nonprivileged user to write a complete process dump and at the same time prevent the user from reading that process dump. To do so, perform the following steps:

1. If the IMGDMPS\$PROTECT rights identifier does not exist, run the image SYS\$SYSTEM:IMGDMP_RIGHTS.EXE to create it (see Section 16.17.1).
2. Create a protected directory with rights identifier IMGDMPS\$PROTECT. For example:

```
$ CREATE /DIRECTORY DKA300:[PROCDUMPS] -
  /PROTECTION=(S:RWE,O:RWE,G,W) /OWNER UIC=IMGDMPS$PROTECT
$ SET SECURITY DKA300:[000000]PROCDUMPS.DIR -
  /ACL=((DEFAULT PROTECTION,SYSTEM:RWED,OWNER:RWED,GROUP:,WORLD:), -
  (IDENTIFIER=IMGDMPS$PROTECT,ACCESS=READ+WRITE), -
  (IDENTIFIER=IMGDMPS$PROTECT,OPTIONS=DEFAULT, -
  ACCESS=READ+WRITE+EXECUTE+DELETE+CONTROL), -
  (CREATOR,ACCESS=NONE))
```

3. Define the executive-mode logical name SYS\$PROTECTED_PROCDMP to point to the protected directory. For example:

```
$ DEFINE /SYSTEM /EXECUTIVE_MODE SYS$PROTECTED_PROCDMP DKA300:[PROCDUMPS]
```

4. If DISKQUOTA is to be used on the disk containing the protected directory, specify the maximum disk space to be used for process dumps. For example:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> DISKQUOTA CREATE /DEVICE=DKA300 ! if necessary
SYSMAN> DISKQUOTA ENABLE /DEVICE=DKA300 ! if necessary
SYSMAN> DISKQUOTA ADD IMGDMPS$PROTECT /DEVICE=DKA300 /PERMQUOTA=10000
SYSMAN> DISKQUOTA REBUILD /DEVICE=DKA300 ! if necessary
SYSMAN> EXIT
```

Warning

Do not grant IMGDMPS\$PROTECT to any user. It is granted and revoked as needed by SYS\$SHARE:IMGDMP.EXE from executive mode while writing a process dump. If you grant it permanently to a user, then that user has access to all process dumps written to the protected directory.

You can choose to set up additional ACLs on the protected directory to further control which users are allowed to read and write process dumps there.

Note that to take a process dump when the image is installed with elevated privileges or belongs to a protected subsystem, the user must hold CMKRNL privilege, and is by definition a privileged user (see Section 16.17.1).

Performance Considerations

This chapter introduces the basic concepts of performance management. For more detailed information, refer to *OpenVMS Performance Management*.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Knowing your work load	Section 17.2
Choosing a work load management strategy	Section 17.3
Distributing the work load	Section 17.4
Predicting when tuning is required	Section 17.6
Evaluating tuning success	Section 17.7
Choosing performance options	Section 17.8
Installing images with the Install utility (INSTALL)	Section 17.10
Reserving memory for specific uses (Alpha only)	Section 17.11

This chapter explains the following concepts:

Concept	Section
Performance management	Section 17.1
System tuning	Section 17.5
Images and known images	Section 17.10.1
Known file lists	Section 17.10.2
Attributes of known images	Section 17.10.3

17.1 Understanding Performance Management

Performance management means optimizing your hardware and software resources for the current work load. This task entails several distinct but related activities:

- Acquiring a thorough familiarity with your work load and an understanding of how that work load exercises the system's resources. This knowledge, combined with an appreciation of the operating system's resource management mechanisms, will enable you to establish realistic standards for system performance in areas such as the following ones:
 - Interactive and batch throughput
 - Interactive response time
 - Batch job turnaround time

Performance Considerations

17.1 Understanding Performance Management

- Routinely monitoring system behavior to determine if, when, and why a given resource is approaching capacity.
- Investigating reports of degraded performance from users.
- Planning for changes in the system work load or hardware configuration and being prepared to make any necessary adjustments to system values.
- Performing, after installation, certain optional system management operations.

17.2 Knowing Your Work Load

One of the most important assets that a system manager brings to any performance evaluation is an understanding of the normal work load and behavior of the system. Each system manager must assume the responsibility for understanding the system's work load sufficiently to be able to recognize normal and abnormal behavior; to predict the effects of changes in applications, operations, or usage; and to recognize typical throughput rates. The system manager should be able to answer such questions as the following ones:

- What is the typical number of users on the system at any given time of day?
- What is the typical response time for various tasks for this number of users, at any given hour of operation?
- What are the peak hours of operation?
- Which jobs typically run at which time of day?
- Which commonly run jobs are intensive consumers of the CPU, memory, and disk space?
- Which applications involve the most image activations?
- Which parts of the system software, if any, have been modified or user-written, such as device drivers?
- Do any known, or anticipated, system bottlenecks exist?

If you are new to the OpenVMS operating system or to system management, you should observe system operation using the following tools:

- Monitor utility
- Accounting utility
- SHOW commands (available through DCL)

The *OpenVMS Performance Management* provides detailed procedures for using the Monitor utility and, to a lesser extent, other operating system tools to observe and evaluate system performance. Also, the *OpenVMS System Management Utilities Reference Manual* provides reference information about using the Monitor utility.

Over time you will learn about metrics such as the typical page fault rate for your system, the typical CPU usage, the normal memory usage, and typical modes of operation. You will begin to see how certain activities affect system performance and how the number of users or the time of day affects some of the values.

As you continue to monitor your system, you will come to know what range of values is acceptable, and you will be better prepared to use these same tools, together with your knowledge, to detect unusual conditions. Routine evaluation of the system is critical for effective performance management. The best way to avoid problems is to anticipate them; you should not wait for problems to develop before you learn how the system performs.

You can learn more about your system's operation if you use the Monitor and Accounting utilities on a regular basis to capture and analyze certain key data items. By observing and collecting this data, you will also be able to see usage trends and predict when your system may reach its capacity.

You should also understand that system resources are used by system management tools. Be careful, therefore, in selecting the items you want to measure and the frequency with which you collect the data. If you use the tools excessively, the consumption of system resources to collect, store, and analyze the data can distort your picture of the system's work load and capacity. The best approach is to have a plan for collecting and analyzing the data.

17.3 Choosing a Work Load Management Strategy

System performance is directly proportional to the efficiency of work load management. Each installation must develop its own strategy for work load management. Before adjusting any system values, make sure you resolve the following issues:

- Does the work load “peak” at a particular time of day, that is, is it noticeably heavier than at other times?
- Can the work load be better balanced? Perhaps some voluntary measures can be adopted by users, after appropriate discussion.
- Could some jobs be run better as batch jobs, preferably during nonpeak hours?
- Have primary and secondary hours of operation been employed with users? If not, could system performance benefit by adopting this practice? If the primary and secondary hours are in use, are the choices of hours the most appropriate for all users? (Plan to review this issue every time you either add or remove users or applications, to ensure that the desired balance is maintained.)
- Can future applications be designed to work around any known or expected system bottlenecks? Can present applications be redesigned somewhat, for the same purpose? (Refer to the *Guide to OpenVMS File Applications*.)
- Are you making the best use of the code-sharing ability that the operating system offers? Code sharing provides an excellent means to conserve memory and improve performance over the life of the system.

17.4 Distributing the Work Load

You should distribute the work load as evenly as possible over the time your system is running. Although the work schedule for your site may make it difficult to schedule interactive users at optimum times, the following techniques may be helpful:

- Run large jobs as batch jobs—Establish a site policy that encourages the submission of large jobs on a batch basis. Regulate the number of batch

Performance Considerations

17.4 Distributing the Work Load

streams so that batch usage is high when interactive usage is low. You might also want to use DCL command qualifiers to run batch jobs at lower priority, adjust the working set sizes, or control the number of concurrent jobs. For information about setting up your batch environment, see Section 14.2.1.

- Restrict system use—Do not permit more users to log in at one time than the system can support with an adequate response time. You can restrict the number of interactive users with the DCL command SET LOGINS/INTERACTIVE. You can also control the number of concurrent processes with the MAXPROCESSCNT system parameter, and the number of remote terminals allowed to access the system at one time with the RJOBLIM system parameter. See Section 15.5 for information about modifying system parameters. Refer to the *OpenVMS System Management Utilities Reference Manual* for descriptions of all system parameters.

You might also restrict use of the system by groups of users to certain days and hours of the day. You can use the Authorize utility to define the permitted login hours for each user. In particular, refer to the AUTHORIZE qualifier /PRIMEDAYS. For more information, see Chapter 7 and the AUTHORIZE section of the *OpenVMS System Management Utilities Reference Manual*.

You can use the DCL command SET DAY to override the conventional day of the week associations for primary and secondary days. For example, you might want to specify a primary day of the week as a secondary day when it is a holiday.

- Design applications to reduce demand on binding resources—If you know where your system bottlenecks are or where they will likely occur in the near future, you can distribute the work load more evenly by planning usage that minimizes demand on any bottleneck points. (Refer to the *Guide to OpenVMS File Applications*.)

17.5 Understanding System Tuning

Tuning is the process of altering various system values to improve *overall* performance possible from any given configuration and work load. However, the process does not include the acquisition and installation of additional memory or devices, although in many cases such additions (when made at the appropriate time) can vastly improve system operation and performance.

On most systems, the work load is constantly changing. System parameters that produce optimal performance at one time may not produce optimal performance a short time later as the work load changes. Your goal is to establish values that produce acceptable performance under all of the changing work load conditions.

Before you undertake any action, you must recognize that the following sources of performance problems cannot be cured by adjusting system values:

- Improper operation
- Unreasonable performance expectations
- Insufficient memory for the applications attempted
- Inadequate hardware configuration for the work load, such as too slow a processor, too few buses for the devices, too few disks, and so forth
- Improper device choices for the work load, such as using disks with insufficient speed or capacity
- Hardware malfunctions

Performance Considerations

17.5 Understanding System Tuning

- Poor application design
- Allowing one process to consume all available resources

When you make adjustments, you normally select a very small number of values for change, based on a careful analysis of the behavior you observed. You control system resources by tuning the values of two types of parameters:

Parameter Type	Description
System parameters	<p>The values set for system parameters control system resources on a systemwide basis. The AUTOGEN command procedure automatically sets system parameters to appropriate values for your system configuration. AUTOGEN can also record feedback from a running system to adjust those parameters based on the system's work load.</p> <p>The <i>OpenVMS Performance Management</i> describes how to select the parameters and new values that are likely to produce the desired changes.</p> <p>Section 15.5 explains how to use AUTOGEN to modify system parameter values.</p>
UAF limits and quotas	<p>The values set for limits and quotas in each user authorization file (UAF) record control system resources on a per-user basis. To control these values, use the Authorize utility. For information, see Section 7.11.</p>

Before you undertake any tuning operation, be sure you are familiar with the resource management mechanisms described in the *OpenVMS Performance Management*. Understand the nature of system values before adjusting them. Without the proper level of understanding, you might degrade, rather than improve, overall performance.

17.6 Predicting When Tuning Is Required

Under most conditions, tuning is rarely required for OpenVMS systems. The AUTOGEN command procedure, which is included in the operating system, establishes initial values for all the configuration-dependent system parameters so that they match your particular configuration. For information about AUTOGEN, see Section 15.4.

Additionally, the system includes features that, in a limited way, permit it to adjust itself dynamically during operation. That is, the system detects the need for adjustment in certain areas, such as the nonpaged dynamic pool, the working set size, and the number of pages on the free and modified page lists. The system makes rough adjustments in these areas automatically. As a result, these areas can grow dynamically, as appropriate, during normal operation.

Experience has shown that the most common cause of disappointment in system performance is insufficient hardware capacity. Once the demand on a system exceeds its capacity, adjusting system values will not result in any significant improvements, simply because such adjustments are a means of trading off or juggling existing resources.

Although tuning is rarely required, you should recognize that system tuning may be needed under the following conditions:

- If you have adjusted your system for optimal performance with current resources and then acquire new capacity, you must plan to compensate for the

Performance Considerations

17.6 Predicting When Tuning Is Required

new configuration. In this situation, the first and most important action is to execute the AUTOGEN command procedure.

- If you anticipate a dramatic change in your work load, you should expect to compensate for the new work load.

17.7 Evaluating Tuning Success

Whenever you adjust your system, you should monitor its behavior afterward to be sure that you have obtained the desired results. To observe results, use the Monitor utility and the various forms of the DCL command SHOW. Refer to the *OpenVMS DCL Dictionary* for detailed information about the SHOW command. See Section 20.8.2 for information about using MONITOR. Refer to the *OpenVMS System Management Utilities Reference Manual* for detailed descriptions of MONITOR commands.

For example, you might consider running some programs (whose results you believe are fixed and reproducible) at the same time that you run your normal work load. If you run the programs and measure their running times under nearly identical work load conditions both before and after your adjustments, you can obtain a basis for comparison.

However, when applying this technique, remember to take the measurements under very similar work load conditions. Also, remember that this test alone does not provide conclusive proof of success. The possibility always exists that your adjustments may have favored the performance of the image you are measuring—to the detriment of other images. Therefore, in all cases, continue to observe system behavior closely for a time after you make any changes.

17.8 Choosing Performance Options

The following list describes optional system management operations normally performed after installation. These operations often result in improved overall performance. Choose the options that are appropriate for your site. Not all options are appropriate at every site.

- Expand system libraries — Most large libraries ship on the operating system in a data-reduced (compressed) format to conserve disk space. Unless the files are expanded (decompressed), the system must dynamically expand them each time they are accessed. The resulting performance slowdown is especially noticeable during link operations and requests to online help. If your system has sufficient disk space, expanding the libraries improves both CPU and elapsed time performance. For details about expanding system libraries and using the LIBDECOMP.COM command procedure, see Section 17.9.
- Disable file system high-water marking—This security feature is set by default when a volume is initialized to guarantee that users cannot read data they have not written.

For nonshared sequential files, the performance impact of high-water marking is minimal. However, for files of nonsequential format, high-water marking creates some overhead; the system erases the previous contents of the disk blocks allocated every time a file is created or extended.

Disabling the feature improves system performance by a variable amount, depending on the following factors:

- How frequently new files are created

Performance Considerations

17.8 Choosing Performance Options

- For indexed and relative files, how frequently existing files are extended
- How fragmented the volume is

Be sure to consider the security implications before you disable high-water marking.

To disable high-water marking, you can specify the /NOHIGHWATER qualifier when initializing the volume, or you can disable high-water marking with the DCL command SET VOLUME in the following format:

```
SET VOLUME/NOHIGHWATER_MARKING device-name[:]
```

- Set file extend parameters for OpenVMS Record Management Services (RMS)—Because files extend in increments of twice the multiblock count (default 16), system defaults provide file extension of 32 blocks rounded up to the nearest multiple of the disk's cluster size. Thus, when files are created or extended, increased I/O may slow performance. The problem can be corrected by specifying larger values for file extend parameters or by setting the system parameter RMS_EXTEND_SIZE. See Section 15.5 for information about modifying system parameters. Refer to the *OpenVMS System Management Utilities Reference Manual* for a description of all system parameters.

For more detailed information about establishing the file extension quantity, refer to the section on tuning in the *Guide to Creating OpenVMS Modular Procedures*.

- Install frequently used images—When an image is accessed concurrently by more than one process on a routine basis, install the image with the Install utility (INSTALL), specifying the /OPEN, /SHARED, and /HEADER_RESIDENT qualifiers. You will thereby ensure that all processes use the same physical copy of the image, and that the image will be activated in the most efficient way.

Generally, an image takes about two additional physical pages when installed with the /OPEN, /HEADER_RESIDENT, and /SHARED qualifiers. The utility's LIST/FULL command shows the highest number of concurrent accesses to an image installed with the /SHARED qualifier. This information can help you decide whether installing an image is an efficient use of memory.

See Section 17.10.11 and the INSTALL section of *OpenVMS System Management Utilities Reference Manual* for more information about installing images.

- On Alpha systems, install shareable and executable images specifying the /RESIDENT qualifier with the Install utility. For more information, see Section 17.10.6.

Note that this is a tradeoff between the CPU and memory. Installing an image with /RESIDENT qualifier means that the code is to be nonpaged. Depending on the amount of sharing, this is can be a memory gain or loss.

- Reduce system disk I/O—You can move frequently accessed files off the system disk and use logical names to specify the location or, where necessary, other pointers to access them. For example:
 - SYSUAF.DAT (SYSUAF is the logical name)
 - RIGHTSLLIST.DAT (RIGHTSLLIST is the logical name)
 - VMSMAIL_PROFILE.DATA (VMSMAIL_PROFILE is the logical name)
 - NETPROXY.DAT (NETPROXY is the logical name)

Performance Considerations

17.8 Choosing Performance Options

- NET\$PROXY.DAT (NET\$PROXY is the logical name)
- The queue database (for more information, see Section 13.3)
- ERRFMT log files (SYS\$ERRORLOG is the logical name)
- MONITOR log files (SYS\$MONITOR is the logical name)
- The accounting log file (ACCOUNTNG is the logical name)
- SECURITY_AUDIT.AUDIT\$JOURNAL (SET AUDIT/JOURNAL=SECURITY/DESTINATION= *filespec*)
- Default DECnet for OpenVMS accounts (records included in the SYSUAF file on the OpenVMS distribution kit)

To redefine logical names for these system files, edit the site-specific command procedure SYS\$MANAGER:SYLOGICALS.COM. For more information about defining logical names in SYLOGICALS.COM, see Section 5.2.5.

You can also consider moving paging and swapping activity off the system disk by creating large secondary page and swap files on a less heavily used disk. However, if you want to store crash dumps for diagnosing system failures, the dump file must reside in the system-specific directory SYS\$SPECIFIC:[SYSEXE] on the system disk for storing crash dumps; if no dump file exists in SYS\$SPECIFIC:[SYSEXE], the primary page file must be located there if you want to store crash dumps. For detailed information about moving page and swap files, see Section 16.16.

- On Alpha systems, set aside large amounts of memory for use within memory-resident sections. The Reserved Memory Registry through its interface within the SYSMAN utility allows an OpenVMS Alpha system to be configured with large amounts of memory set aside for use within memory-resident sections and by other privileged applications. The Reserved Memory Registry also allows an OpenVMS system to be properly tuned through the AUTOGEN utility, taking into account the preallocated reserved memory. For more information, see Section 17.11.

17.9 Expanding System Libraries

Some of the larger system libraries are shipped with the OpenVMS operating system in a data-reduced (compressed) format. Expanding (that is, decompressing) these libraries gives the system faster access to them but also consumes more disk space. To learn how to determine whether you have enough disk space to expand some or all libraries, refer to Section 17.9.1.

If you have delete access to the existing libraries and write access to the library directories, you can use either of the following methods to expand and reduce (compress) libraries:

- Run the LIBDECOMP.COM command procedure (see Section 17.9.2), which can expand or reduce one, some, or all of the selected libraries that ship in data-reduced (compressed) format.
- Enter the DCL command LIBRARY with the /DATA qualifier (see Section 17.9.3) to expand or reduce one library at a time, including libraries that do not ship in data-reduced format.

Note

In this discussion, “decompress” and “compress” refer to actions that result from using either the LIBDECOMP.COM procedure or the LIBRARY command with the /DATA=EXPAND or /DATA=REDUCE qualifier. These actions should not be confused with the results of specifying the /COMPRESS qualifier, which performs a different, unrelated function. For more information about the LIBRARY command and its qualifiers, refer to the *OpenVMS Command Definition, Librarian, and Message Utilities Manual* or online help.

17.9.1 Determining Disk Space Available to Expand Libraries

Before you expand any libraries, make sure your system has enough free disk space to accommodate the expanded files. To find out how much free disk space you have, enter the following command:

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

The amount of space required to expand all the libraries that ship in data-reduced format differs, depending on your operating system and configuration. For Alpha systems, you need approximately 60,000 free disk blocks; VAX systems require less. For specific disk requirements, refer to the *Upgrade and Installation Manual* for your operating system.

If you have less free disk space than is required, or if you do not want to expand all the system libraries, you can choose to expand only some of them. For example, you might expand only the system help library (HELPLIB.HLB) and other libraries that are used frequently on your system, and leave the rest in data-reduced format.

17.9.2 Using the Library Decompression Utility (LIBDECOMP.COM)

The Library Decompression utility is a command procedure, LIBDECOMP.COM, located in the SYS\$UPDATE directory. LIBDECOMP.COM uses LIBRARY commands to expand (decompress) or reduce (compress) any or all of the system libraries that ship in data-reduced format. (For a list of these libraries, see Section 17.9.2.1.)

The LIBDECOMP.COM command procedure allows you to enter one command instead of multiple LIBRARY commands to expand or reduce all or many libraries in one operation. LIBDECOMP.COM can be executed interactively (see Section 17.9.2.2) or in batch mode (see Section 17.9.2.3).

17.9.2.1 Libraries on which LIBDECOMP.COM Operates

The Library Decompression utility works on a selected set of large libraries — those that ship in data-reduced format. Some libraries are platform specific and ship only on VAX systems or only on Alpha systems. Many libraries are optional components and may not be present on your system if all options were not selected when OpenVMS was installed. Table 17–1 lists all the libraries known to LIBDECOMP.COM on both VAX and Alpha systems.

Performance Considerations

17.9 Expanding System Libraries

Table 17–1 Libraries Recognized by LIBDECOMP.COM

Library Name	Platform	Description
[SYSHLP] directory; help library files (.HLB)		
ACLEDT.HLB	Both	Access Control List Editor help
BKM\$HELP.HLB	Both	Backup Manager help
DBG\$HELP.HLB	Both	OpenVMS Debugger help
DBG\$UIHELP.HLB	Both	OpenVMS Debugger help
EDTHELP.HLB	Both	EDT editor help
EVE\$HELP.HLB	Both	EVE editor help
EVE\$KEYHELP.HLB	Both	EVE keypad help
EXCHNGHLP.HLB	Both	Exchange utility help
HELPLIB.HLB	Both	DCL help
LANCP\$HELP.HLB	Both	LAN Control Program help
LATCP\$HELP.HLB	Both	LAT Control Program help
MAILHELP.HLB	Both	Mail utility help
NCPHELP.HLB	Both	Network Control Program help
SDA.HLB	Both	System Dump Analyzer help
SHWCLHELP.HLB	Both	Show Cluster utility help
SYSGEN.HLB	Both	System Generation utility help
SYSMANHELP.HLB	Both	System Management utility help
TPUHELP.HLB	Both	Text Processing Utility help
UAFHELP.HLB	Both	Authorize utility help
[SYSLIB] directory; macro library files (.MLB)		
LANIDF.MLB	Alpha only	LAN internal driver macros
LIB.MLB	Both	Operating system macros
STARLET.MLB	Both	Operating system macros
SYSLDMLB.MLB	VAX only	System build files

(continued on next page)

Performance Considerations 17.9 Expanding System Libraries

Table 17–1 (Cont.) Libraries Recognized by LIBDECOMP.COM

Library Name	Platform	Description
[SYSLIB] directory; object library files (.OLB)		
DECCRTL.OLB	VAX only	Compaq C Run-time Library
STARLET.OLB	Both	System object library and run-time library
SYSBLDLIB.OLB	VAX only	System build files
VAXCRTL.OLB	Both	Compaq C RTL routine name entry points; VAX G_floating double-precision, floating-point entry points
VAXCRTLD.OLB	Alpha only	Limited support of VAX D_floating double-precision, floating-point entry points
VAXCRTLDX.OLB	Alpha only	VAX D_floating support; support for /L_DOUBLE_SIZE=128 compiler qualifier
VAXCRTLT.OLB	Alpha only	IEEE T_floating double-precision, floating-point entry points
VAXCRTLTX.OLB	Alpha only	IEEE T_floating support; support for /L_DOUBLE_SIZE=128 compiler qualifier
VAXCRTLX.OLB	Alpha only	G_floating support; support for /L_DOUBLE_SIZE=128 compiler qualifier
VMS\$VOLATILE_PRIVATE_INTERFACES.OLB	Alpha only	OpenVMS bugcheck processing codes
[SYSLIB] directory; text library files (.TLB)		
ERFLIB.TLB	Both	ANALYZE/ERROR device descriptions
LIB_ADA_SUBSET.TLB	Both	Ada programmers toolkit of operating system definitions
NTA.TLB	Both	Files to build against NTA facility
STARLET_RECENT_ADA_SUBSET.TLB	Both	Ada programmers toolkit of operating system definitions
STARLETSD.TLB	Both	STARLET definitions used during layered product installations
SYS\$LIB_C.TLB	Alpha only	Header files for C language; derived from LIB
SYS\$STARLET_C.TLB	Both	Public header files for Compaq C

You can use the list function of the LIBDECOMP.COM command procedure to output a list of all libraries known to the Library Decompression utility, including their size and status on your system. For details and an example, see Section 17.9.2.2.1.

17.9.2.2 Using LIBDECOMP.COM Interactively

The basic command to run the Library Decompression utility is as follows:

```
@SYS$UPDATE:LIBDECOMP [parameters]
```

The Library Decompression utility accepts up to eight optional parameters. The first parameter controls which of the utility's three functions is performed. Additional parameters control which libraries the utility processes. The three functions are described in the following sections:

- List function (see Section 17.9.2.2.1)

Performance Considerations

17.9 Expanding System Libraries

Lists all VAX and Alpha libraries known to the Library Decompression utility and displays the size and status (reduced or expanded) of libraries on your system.

- Expand function (see Section 17.9.2.2.2)
Expands libraries that are in data-reduced format. This is the default function.
- Reduce function (see Section 17.9.2.2.3)
Compresses expanded libraries into data-reduced format.

To get a brief online help display, enter the following command:

```
$ @SYS$UPDATE:LIBDECOMP HELP
```

17.9.2.2.1 Listing Libraries To list all VAX and Alpha libraries known to the Library Decompression utility, along with their size and status on your system, specify the LIST parameter in the command, as follows:

```
$ @SYS$UPDATE:LIBDECOMP LIST
```

The resulting list indicates which libraries are not present on your system, either because they do not ship with VAX or with Alpha or because the facility associated with them is not installed on your system. For libraries that are present, the list includes the library size and the current status (reduced or expanded). The following example shows the output from an Alpha system. (Note: File sizes are subject to change. For the most accurate information, refer to the LIST output on your own system.)

Performance Considerations 17.9 Expanding System Libraries

\$ @SYS\$UPDATE:LIBDECOMP LIST

OpenVMS Library Decompression Utility
List of all libraries known to LIBDECOMP
"Library not present" indicates not installed on this system

Libraries in SYS\$SYSROOT:

Library	Size	
1) [SYSHLP]ACLEDT.HLB	70	Reduced format
2) [SYSHLP]BKM\$HELP.HLB	156	Reduced format
3) [SYSHLP]DBG\$HELP.HLB	1234	Reduced format
4) [SYSHLP]DBG\$UIHELP.HLB	269	Reduced format
5) [SYSHLP]EDTHELP.HLB	154	Reduced format
6) [SYSHLP]EVE\$HELP.HLB	676	Reduced format
7) [SYSHLP]EVE\$KEYHELP.HLB	99	Reduced format
8) [SYSHLP]EXCHNGHLP.HLB	83	Reduced format
9) [SYSHLP]HELPLIB.HLB	9179	Reduced format
10) [SYSHLP]LANCP\$HELP.HLB	119	Reduced format
11) [SYSHLP]LATCP\$HELP.HLB	157	Reduced format
12) [SYSHLP]MAILHELP.HLB	211	Reduced format
13) [SYSHLP]NCPHELP.HLB	261	Reduced format
14) [SYSHLP]SDA.HLB	308	Reduced format
15) [SYSHLP]SHWCLHELP.HLB	103	Reduced format
16) [SYSHLP]SYSGEN.HLB	337	Reduced format
17) [SYSHLP]SYSMANHELP.HLB	492	Reduced format
18) [SYSHLP]TPUHELP.HLB	575	Reduced format
19) [SYSHLP]UAFHELP.HLB	241	Reduced format
20) [SYSLIB]LANIDEF.MLB	181	Reduced format
21) [SYSLIB]LIB.MLB	2715	Reduced format
22) [SYSLIB]STARLET.MLB	2335	Reduced format
23) [SYSLIB]SYSBLDMLB.MLB		Library not present
24) [SYSLIB]DECCRTL.OLB		Library not present
25) [SYSLIB]STARLET.OLB	27461	Reduced format
26) [SYSLIB]SYSBLDLIB.OLB		Library not present
27) [SYSLIB]VAXCRTL.OLB	1163	Reduced format
28) [SYSLIB]VAXCRTLD.OLB	1587	Reduced format
29) [SYSLIB]VAXCRTLDX.OLB	1506	Reduced format
30) [SYSLIB]VAXCRTLT.OLB	1434	Reduced format
31) [SYSLIB]VAXCRTLTX.OLB	1449	Reduced format
32) [SYSLIB]VAXCRTLX.OLB	1285	Reduced format
33) [SYSLIB]ERFLIB.TLB	64	Reduced format
34) [SYSLIB]LIB_ADA_SUBSET.TLB	1839	Reduced format
35) [SYSLIB]NTA.TLB	34	Reduced format
36) [SYSLIB]STARLETSD.TLB	3940	Reduced format
37) [SYSLIB]SYS\$LIB_C.TLB	9442	Reduced format
38) [SYSLIB]SYS\$STARLET_C.TLB	5864	Reduced format
39) [SYSLIB]VMS\$VOLATILE_PRIVATE_INTERFACES.OLB	445	Reduced format
40) [SYSLIB]STARLET_RECENT_ADA_SUBSET.TLB	1100	Reduced format
Total Libraries: 37		78568

17.9.2.2 Expanding (Decompressing) Libraries The default action for LIBDECOMP.COM is the expand function. You can explicitly specify EXPAND as the first parameter on your command line, but this is not necessary. Unless the first parameter is LIST or REDUCE, the expand function is used by default.

If the expand function is used, the remaining parameters specify which of the libraries to expand. You can specify ALL to expand all libraries known to the utility, or you can specify up to eight library names (seven names if EXPAND is specified). Wildcard characters are not permitted. The specified libraries must be

Performance Considerations

17.9 Expanding System Libraries

known to the utility. (To expand any other libraries, you must use the LIBRARY command, as described in Section 17.9.3.) If you do not specify ALL or a list of libraries, LIBDECOMP.COM prompts you for the libraries to expand. You can specify any number of libraries in response to the prompt.

Note

Expanding or reducing all of the libraries known to LIBDECOMP.COM will generally take about five to ten minutes. However, depending on the specific hardware and software configuration of your system, and on other activity, this can take longer, up to half an hour or more in some cases.

The expand function produces a display similar to the one shown in the following OpenVMS Alpha example. After the header lines are displayed, there is a pause while LIBDECOMP.COM checks the status of each library.

```
OpenVMS Library Decompression Utility
Candidate Libraries to be expanded
(Libraries not present and libraries already expanded are not listed)

  1 ACLEDT.HLB          13 NCPHELP.HLB          25 VAXCRTLD.OLB
  2 BKM$HELP.HLB       14 SDA.HLB              26 VAXCRTLDX.OLB
  3 DBG$HELP.HLB       15 SHWCLHELP.HLB       27 VAXCRTLT.OLB
  4 DBG$UIHELP.HLB    16 SYSGEN.HLB          28 VAXCRTLT_X.OLB
  5 EDTHELP.HLB        17 SYSMANHELP.HLB     29 VAXCRTLX.OLB
  6 EVE$HELP.HLB       18 TPUHELP.HLB         30 ERFLIB.TLB
  7 EVE$KEYHELP.HLB   19 UAFHELP.HLB         31 LIB_ADA_SUBSET.TLB
  8 EXCHNGHLP.HLB     20 LANIDF.MLB          32 NTA.TLB
  9 HELPLIB.HLB        21 LIB.MLB             33 STARLETSD.TLB
 10 LANCP$HELP.HLB    22 STARLET.MLB         34 SYS$LIB_C.TLB
 11 LATCP$HELP.HLB    23 STARLET.OLB         35 SYS$STARLET_C.TLB
 12 MAILHELP.HLB      24 VAXCRTL.OLB

 36 VMS$VOLATILE_PRIVATE_INTERFACES.OLB
 37 STARLET_RECENT_ADA_SUBSET.TLB

A ALL libraries to be expanded
H Display HELP information for LIBDECOMP
E EXIT this procedure
```

If you specified ALL, the following message is displayed, and the utility then expands all listed libraries.

```
"ALL" specified; all libraries will be processed
```

If you did not specify ALL, you are prompted as follows:

```
* Enter a letter or the number(s) of libraries to be expanded
  (Separate multiple numbers with a comma)
```

Enter A, H, E, or one or more numbers to indicate which libraries to expand. There is no limit to how many numbers you can specify.

If you enter parameters to identify specific libraries to expand, LIBDECOMP.COM does not display the list as shown in the example; rather, it lists each library as it is processed.

Performance Considerations

17.9 Expanding System Libraries

Examples

- Either of the following commands expands all libraries:

```
$ @SYS$UPDATE:LIBDECOMP ALL
$ @SYS$UPDATE:LIBDECOMP EXPAND ALL
```

- To choose whether to expand some or all of the libraries in a menu, first enter one of the following commands:

```
$ @SYS$UPDATE:LIBDECOMP
$ @SYS$UPDATE:LIBDECOMP EXPAND
```

Then, as shown in an earlier example, you are prompted to specify all libraries or to specify the numbers for an unlimited number of individual libraries to be expanded.

- To bypass the menu and expand only selected libraries, enter a command similar to one of the following, in which you specify the library names in the command line instead of responding to a prompt:

```
$ @SYS$UPDATE:LIBDECOMP HELPLIB.HLB STARLET.MLB LIB.MLB
$ @SYS$UPDATE:LIBDECOMP EXPAND HELPLIB STARLET.MLB LIB
```

Both commands have the same result: to expand the HELPLIB.HLB, STARLET.MLB, and LIB.MLB libraries.

17.9.2.2.3 Reducing (Compressing) Libraries When you specify REDUCE as the first parameter in your command line, LIBDECOMP.COM reduces libraries that have been expanded.

You can specify ALL after REDUCE to reduce all of the libraries known to the utility, or you can specify a list of up to seven libraries. Wildcard characters are not permitted.

If you do not specify ALL or the name of at least one library, LIBDECOMP.COM prompts you for the libraries to reduce. There is no limit to the number of libraries you can list in response to the prompt.

The reduce function produces a display similar to that of the expand function except that it displays only the libraries eligible to be reduced.

Examples

- The following command reduces all libraries that have been expanded:

```
$ @SYS$UPDATE:LIBDECOMP REDUCE ALL
```

- To choose whether to reduce some or all of the libraries in a menu, first enter the following command:

```
$ @SYS$UPDATE:LIBDECOMP REDUCE
```

Then, as with the expand function, you are prompted to specify all libraries or to specify the numbers for an unlimited number of individual libraries to be reduced.

Performance Considerations

17.9 Expanding System Libraries

17.9.2.3 Using LIBDECOMP.COM in Batch Mode

You can submit the Library Decompression utility to a batch queue by using the DCL command `SUBMIT` with the `/PARAMETERS` qualifier, as follows:

```
SUBMIT SYS$UPDATE:LIBDECOMP /PARAMETERS=(p1[,p2,...])
```

The batch procedure produces the same results as the interactive procedure, but with the batch job you must specify `HELP`, `LIST`, `ALL`, or at least one library name, because the batch job cannot prompt you for input.

You can specify up to eight parameters. If you specify more than one parameter, the parameters must be enclosed in parentheses and must be separated by commas.

Examples

- The following command outputs a list of all the libraries known to `LIBDECOMP.COM`:

```
$ SUBMIT SYS$UPDATE:LIBDECOMP /PARAMETERS=LIST
```

- By default, the following command expands all eligible libraries:

```
$ SUBMIT SYS$UPDATE:LIBDECOMP /PARAMETERS=ALL
```

- The following command expands the `HELPLIB.HLB`, `STARLET.MLB`, and `LIB.MLB` libraries:

```
$ SUBMIT SYS$UPDATE:LIBDECOMP -  
_ $ /PARAMETERS=(EXPAND,HELPLIB.HLB,STARLET.MLB,LIB.MLB)
```

17.9.3 Using the LIBRARY Command with the /DATA Qualifier

An alternative to using the Library Decompression utility to expand or reduce an individual library is to use the DCL command `LIBRARY` with the `/DATA` qualifier. With this method, you can specify only one library per `LIBRARY` command.

Unlike `LIBDECOMP.COM`, which acts on only about 40 selected libraries that ship in data-reduced format, the `LIBRARY` command can be used to expand or reduce almost *any* library file, provided the library type (the file extension) is known to the OpenVMS Librarian utility. System libraries that you should not compress include the following:

- `[SYSLIB]EPC$FACILITY.TLB`

This file ships empty.

- `[SYSLIB]IMAGELIB.OLB`

Compression is not effective on this library.

- `[SYSLIB]NCS$LIBRARY.NLB`

Library type `.NLB` is not known to the OpenVMS Librarian utility.

To expand a specified library, use the following command format:

```
LIBRARY library-name.ext /DATA=EXPAND
```

To reduce a specified library, use the following command format:

```
LIBRARY library-name.ext /DATA=REDUCE
```

You must always specify the library extension (`.HLB`, `.MLB`, `.OLB`, or `.TLB`).

If the specified library is not located on your current default device and directory, you must also specify the device and directory in the library specification. Most system libraries are in either SYS\$HELP ([SYSHLP]) or SYS\$LIBRARY ([SYSLIB]). For example:

```
$ LIBRARY [SYSHLP]HELPLIB.HLB /DATA=EXPAND
```

For information about other qualifiers for the LIBRARY command, see online help for LIBRARY or refer to the *OpenVMS Command Definition, Librarian, and Message Utilities Manual*.

17.10 Using INSTALL to Install Known Images

The Install utility (INSTALL) stores information about images in memory. Use INSTALL for the following reasons:

Reason	For More Information
To conserve memory use for images that are used concurrently	Section 17.10.7
To improve system performance	Section 17.10.5
‡On Alpha systems, to improve performance by using images with shared address data	Section 17.10.6
To make executable images that require enhanced privileges available for general use	Section 17.10.8.1
To allow a nonprivileged image to call the privileged functions of a shareable image	Section 17.10.8.2
To mark a sharable image as trusted so it can be invoked by privileged executable images	Section 17.10.9
‡Alpha specific	

The site-independent startup command procedure, STARTUP.COM, uses INSTALL to install certain system images when the system boots. You use INSTALL to install other selected images, according to the needs of your site.

Installed images must be reinstalled each time the system reboots. To do so, include INSTALL commands in the site-specific startup command procedure SYSTARTUP_VMS.COM, as explained in Section 5.2.7.

The Install utility (INSTALL) only installs images that are linked with the /NOTRACEBACK qualifier.

Note that INSTALL commands perform a different function than System Generation utility (SYSGEN) INSTALL commands.

The following sections explain installed images and how to use the Install utility.

17.10.1 Understanding Images and Known Images

An **image** is a collection of procedures and data bound together by the Linker utility. Executable images can be executed (or run) in a process, either by a command line interpreter (CLI) or the \$CREPRC system service. Usually, executable programs have the file type .EXE.

There are three types of images:

Performance Considerations

17.10 Using INSTALL to Install Known Images

Image Type	Description
Executable	An image linked with the /EXECUTABLE qualifier (or without the /SHAREABLE qualifier) of the Linker utility. For more information, refer to the <i>OpenVMS Linker Utility Manual</i> .
Shareable	An image linked with the /SHAREABLE qualifier of the Linker utility. Shareable images are sometimes referred to as linkable images because they can be specified—implicitly or explicitly—as input files to the link of another file. A shareable image is not copied into the executable images that link with it. Thus, only one copy of the shareable image needs to be on disk, no matter how many executable images have linked with it. For more information, refer to the <i>OpenVMS Linker Utility Manual</i> .
System	An image that does not run under the control of the operating system. It is intended for standalone operation only. The content and format of a system image differs from that of shareable images and executable images. For more information, refer to the <i>OpenVMS Linker Utility Manual</i> .

When you install an image with `INSTALL`, the image is assigned attributes and becomes known to the system. For this reason, an installed image is also called a **known image**.

The image activator processes search lists in two passes, in order to favor known images. On its first pass through the search list, the image activator looks up images as known files. If needed, on a second pass through the search list, the image activator looks up images on disk.

17.10.2 Understanding Known File Entries

The system defines known images in internal data structures called **known file entries**. Each entry identifies the file name of the installed image and the attributes with which it was installed (for information about attributes of installed images, see Section 17.10.3).

Known file entries last only while the system is operating. If the system is shut down or fails for any reason, you must reinstall all known images after the system is rebooted.

17.10.3 Understanding Attributes You Can Assign to Known Images

By specifying appropriate qualifiers to `INSTALL` commands, you can assign attributes to known images. Table 17–2 describes these attributes and the qualifiers that are used to assign them to known images.

Table 17–2 Attributes of Known Images

Attribute	Description	Qualifier
Header resident	The header of the image file (native images only) remains permanently resident, saving one disk I/O operation per file access. For images with single-block file headers, the cost is less than 512 bytes of paged dynamic memory per file; for images with multiblock headers, the cost varies according to the header block count. Images installed header resident are implicitly installed permanently open.	/[NO]HEADER_RESIDENT

(continued on next page)

Table 17–2 (Cont.) Attributes of Known Images

Attribute	Description	Qualifier
Permanently open	The image file remains open, so access to the image does not require a call to the file system.	/OPEN
Privileged	Amplified privileges are temporarily assigned to any process running the image, permitting the process to exceed its user authorization file (UAF) privilege restrictions during execution of the image. In this way, users with normal privileges can run programs that require higher-than-normal privileges. This attribute (and the /PRIVILEGED qualifier that creates it) applies only to executable images.	/PRIVILEGED[=(priv-name[,...])]
Protected	When the image is activated, the address space for the image is protected against modification by user-mode code. This is critical for shareable code that runs in kernel or executive mode.	/PROTECTED
‡ Resident	On Alpha systems, code or read-only data for an image is made permanently resident in a system region of memory. This improves performance by using a special page mapping to reduce translation buffer (TB) miss rates. The resident attribute applies to shareable or executable images that have been linked with the qualifier /SECTION_BINDING=(CODE,DATA).	/RESIDENT
Shared	More than one user can access the read-only and non-copy-on-reference read/write sections of the image concurrently, so that only one copy of those sections needs to be in physical memory. (Copy-on-reference sections always require a separate copy for each process.) The image is implicitly declared permanently open.	/SHARED
Writable	When a shareable non-copy-on-reference writable section is removed from physical memory (for paging reasons or because no processes are referencing it), it is written back to the image file. Any updates made by processes mapped to the section, therefore, are preserved (while the initial values are lost). The image must also be declared shareable.	/WRITABLE

‡Alpha specific

17.10.4 Determining Which Images to Install

You can install images for the following reasons:

- Improve image activation performance
- On Alpha systems, improve execution performance of shared images
- Conserve physical memory
- Enhance privileges of images during execution

Because an installed file requires system resources, such as paged dynamic memory, install those files that most improve system performance and site requirements. The INSTALL command LIST provides information about installed images to help you evaluate the merits of installing images. For example, the LIST command calculates the number of times each image is accessed, and shows the number of concurrent accesses, so you can determine if the installation of the images is worth the overhead.

Performance Considerations

17.10 Using INSTALL to Install Known Images

17.10.5 Installing Images to Improve Image Activation Performance

You can improve image activation performance by installing images that run frequently. Image activation performance improves when programs are installed because the operating system opens installed files by file ID rather than by file name, thus eliminating costly directory operations.

Installing images as header resident further enhances activation performance because the system avoids the overhead of I/O operations to read the image header into memory.

To install an image as header resident, specify the `/HEADER_RESIDENT` qualifier when you install the image. This makes the header of the image file remain permanently resident, saving disk I/O. Specifying the `/HEADER_RESIDENT` qualifier implicitly makes the images permanently open.

Image headers are stored in paged dynamic memory. The size of the image headers varies.

Frequently accessed images, critical to a site's operations, can be installed as open images. To install an image as permanently open, specify the `/OPEN` qualifier when you install the image. The image file remains open, so access to the image does not require a call to the file system. The cost of keeping an image file permanently open is approximately 512 bytes of nonpaged dynamic memory per file.

17.10.6 Installing Images with Shared Address Data

Using shared address data on OpenVMS Alpha systems improves performance at the following times:

- At run time, shared address data saves physical memory because of increased memory sharing between processes.
- At image activation, shared address data reduces CPU and I/O time because fixup is performed at installation time.

For details, refer to the `INSTALL` section of the *OpenVMS System Management Utilities Reference Manual*.

Related Terms

Explanations of terms related to shared address data follow.

- Image section

An image consists of a number of image sections. An image section can contain:

- Instructions (code)
- Read-only data (constants)
- Read/write data

- Shared known images

You can make an image a shared known image by using the following command:

```
$ INSTALL ADD image-name /SHARED
```

When you enter this command, the Install utility creates global sections for read-only image sections, allowing the sections to be shared by all the processes that run the image.

Performance Considerations

17.10 Using INSTALL to Install Known Images

- **Address data**

One kind of image section contains address data. At execution time, address data sections are read-only. However, the addresses are not known until image activation; therefore, the image section is read/write until the end of image activation. Addresses for a shareable image generally vary with the process because different processes are likely to have different collections of mapped images.
- **Shared address data**

The shared address data feature assigns unique P1 space addresses for shareable images from the P1 image region. (The IMGREG_PAGES system parameter determines the size of the P1 space.) With the assigned address, the Install utility determines the content of an address data section when the image is installed. A global section is created to allow shared access to each address data image section.

Executable (main) images can also use shared address data sections; these images are not assigned P1 addresses, however, because the base address for an executable image is determined when the image is linked.

17.10.6.1 System-Provided Images

Many images that are part of the OpenVMS software product are installed as shared known images with shared address data. This provides the performance benefit without requiring the system manager to take any explicit action.

17.10.6.2 Application Images

As system manager, you might choose to install additional images with shared address data. In considering this option, you need to investigate application dependencies on sharable images.

17.10.7 Installing Images to Conserve Physical Memory

You can conserve physical memory by installing images that usually run concurrently from several processes. When an image is not installed, or is installed without the shared attribute, each process running the image requires private sections in memory. Shared images conserve physical memory because only one copy of the code needs to be in memory at any time, and many users can access the code concurrently. Use the /SHARED qualifier to install images as shared images.

When you install an image with the shared attribute, permanent system global sections are created. Execution of non-copy-on-reference global sections requires only one copy per section to be in physical memory, no matter how many processes are running the image to which the sections belong.

The number of images you can install with the shared attribute is restricted by the GBLPAGES and GBLSECTIONS system parameters. For more information about these system parameters, refer to the *OpenVMS System Management Utilities Reference Manual*.

Performance Considerations

17.10 Using INSTALL to Install Known Images

17.10.8 Installing Images to Enhance Privileges of Images

There are two ways to allow an image to execute in an enhanced privilege environment:

- Installing executable images with privileges to allow a nonprivileged process to perform the privileged functions of the image.

Use the /PRIVILEGED qualifier with the INSTALL commands CREATE or REPLACE.

Caution

Installing an image with enhanced privilege can compromise system security. Make sure the image does not enable a user to regain control with extra privileges enabled.

- Installing protected shareable images (which are used to implement user-written system services), allowing other, nonprivileged images to execute select portions of privileged code without enhancing the privileges of those individual images.

Use the /PROTECTED and /SHARED qualifiers with the INSTALL commands CREATE or REPLACE.

17.10.8.1 Privileged Executable Images

A nonprivileged process can perform the privileged functions of an executable image when the image is installed with privileges. Install executable images with enhanced privileges by using the /PRIVILEGED qualifier; amplified privileges are temporarily assigned to any process running the image (executable images only), permitting the process to exceed its user authorization file (UAF) privilege restrictions during execution of the image. In this way, users with normal privileges can run programs that require higher-than-normal privileges.

For an image installed with privileges to activate another image, such as a shareable image, either by having it linked to the privileged image or by using LIB\$FIND_IMAGE_SYMBOL, the following conditions hold:

- The shareable image must be installed as a known image using INSTALL.
- Logical names and table names used to find the image must be defined in executive or kernel mode. In particular, the standard executive-mode definition of LNM\$FILE_DEV translates only to LNM\$SYSTEM; definitions in the process, job, or group tables are not recognized.
- Only images linked with the Linker utility qualifiers /NODEBUG and /NOTRACE can be installed with enhanced privilege.

17.10.8.2 Privileged Shareable Images

A privileged shareable image is a shareable image with defined entry points that execute in inner (executive or kernel) mode. Inner-mode entry points in shareable images are referred to as user-written system services.

To create a privileged shareable image, you must:

1. Create a program section with the VECTOR attribute containing a PLV (privileged library vector) data structure.

Performance Considerations

17.10 Using INSTALL to Install Known Images

2. Link the shareable image with the /PROTECT command qualifier or the PROTECT= option of the Linker utility, so that the image acquires its particular form of enhanced privileges.
 - Use the /PROTECT command qualifier when all parts of an image require protection.
 - Use the PROTECT= option when only part of a privileged shareable image requires protection.
3. Install the privileged shareable image with the Install utility, specifying both the /PROTECTED and the /SHARED qualifiers. The /PROTECTED qualifier assigns the protected attribute. The /SHARED qualifier assigns the shareable attribute. See Section 17.10.3 for information about these attributes.

Note

You cannot grant privileges to a shareable image using the /PRIVILEGED qualifier for the INSTALL commands ADD or CREATE. This qualifier works only for executable images.

For more information about creating privileged shareable images, refer to the *OpenVMS Programming Concepts Manual*.

17.10.9 Activating Images in a Privileged Context

When a process performs one of the following actions, the image activator enters a restricted mode of operation similar to that entered when a privileged program is run:

- Runs an executable or shareable image to which it has execute but not read access.
- Activates an executable image installed with the /PRIVILEGE or /EXECUTE_ONLY qualifier.
- Activates an image called by a privileged shareable image.

In this mode of operation:

- All shareable images activated during the life of the execute-only image must be installed.
- The image activator directs OpenVMS RMS to use only **trusted logical names** (logical names associated with executive or kernel mode) when opening image files.

Note

The executable image that calls an execute-only shareable image must be installed with the /EXECUTE_ONLY qualifier, which enables the executable image to activate shareable images to which the process has execute but not read access.

The /EXECUTE_ONLY qualifier has meaning only for executable images. This restriction assures that shareable images running in a privileged context can be trusted to behave as expected.

Performance Considerations

17.10 Using INSTALL to Install Known Images

17.10.10 Specifying File Names in INSTALL

When you use INSTALL commands, your file specifications must name existing executable or shareable images. OpenVMS RMS resolves each file specification using the following defaults:

- A device and directory type of SYS\$SYSTEM
- A file type of .EXE

You can specify a specific version of the file as the known version of the image with the CREATE or REPLACE command. Even if other versions of the file exist, the version that you specify will be the version that satisfies all known file lookups for the image.

17.10.11 Installing Images with INSTALL

Before performing this task, you should understand the following points:

- Attributes of installed images. For information, see Section 17.10.3.
- File specifications for the Install utility. For information, see Section 17.10.10.

How to Perform This Task

1. Give yourself the CMKRNL privilege by entering the following command:

```
$ SET PROCESS/PRIVILEGES=CMKRNL
```

2. Invoke INSTALL by entering the following command:

```
$ INSTALL
```

3. Enter the CREATE command in the following format:

```
CREATE file-spec [/qualifier...]
```

Specify one or more of the following qualifiers, depending on which attributes you want to assign to the image:

```
/EXECUTE_ONLY  
/HEADER_RESIDENT  
/OPEN  
/PRIVILEGED  
/PROTECTED  
/RESIDENT (Alpha systems only)  
/SHARED  
/WRITABLE
```

For more information about installing images, refer to the INSTALL command CREATE in the *OpenVMS System Management Utilities Reference Manual*.

Note

Installing the Install utility itself requires that a number of shareable images have been previously installed. If any of those required shareable images (such as SMG\$SHR, LIBOTS, and so on) is unavailable, the execution of the Install utility fails. Since INSTALL will not work in this situation, you cannot simply install the missing images. To work around this problem, redefine the INSTALL command as follows:

Performance Considerations

17.10 Using INSTALL to Install Known Images

```
$ DEFINE INSTALL SYS$SYSTEM:INSTALL.EXE;0
```

When you now enter the INSTALL command, the image activator does not check the known files list for INSTALL.EXE, and the INSTALL command will complete, allowing you to install the required shareable images.

17.10.12 Displaying Known Images with INSTALL

Use the INSTALL command LIST to display information about known images.

The information displayed with the /FULL qualifier of the LIST command can help you determine if installing an image is worth the expense.

How to Perform This Task

1. Invoke INSTALL by entering the following command:

```
$ INSTALL
```

2. To display a list of all known images and their attributes, enter the LIST command. To display attributes for a specific image, specify the name of the image as follows:

```
LIST file-spec
```

For example:

```
INSTALL> LIST LOGINOUT
```

3. To display complete information about a specific image, including the number of accesses, the number of concurrent accesses, and the number of global sections created, specify the /FULL qualifier as follows:

```
LIST/FULL file-spec
```

Example

The following example displays complete information about the installed image LOGINOUT.EXE, including the number of accesses, the number of concurrent accesses, and the number of global sections created:

```
$ INSTALL
INSTALL> LIST/FULL LOGINOUT

DISK$VMS551:<SYS2.SYSCOMMON.SYSEXE>.EXE
  LOGINOUT;2      Open Hdr Shar Prv
    Entry access count      = 36366
    Current / Maximum shared = 1 / 10
    Global section count    = 3
    Privileges = CMKRNL SYSNAM LOG_IO ALTPRI TMPMBX SYSPRV
INSTALL>
```

17.10.13 Defining Logical Names for Shareable Image Files

If a shareable image is not located in SYS\$SHARE, you must define a logical name for that image in order to run an executable image linked against it. For example, if the file specification for STATSHR is SYS\$SHARE:STATSHR.EXE, no logical name is necessary. But if you put STATSHR in SYS\$DEVICE:[TEST], you must define STATSHR as a logical name before running an executable image that calls it. The logical name must be the same one that was used as the input file specification for the shareable image when it was linked (this is the same name used in installation). For example:

```
$ DEFINE STATSHR SYS$SYSDEVICE:[TEST]STATSHR
```

Performance Considerations

17.10 Using INSTALL to Install Known Images

By redefining the logical name of a shareable image, you can replace that shareable image with another without requiring the calling executable image to relink. For example, the following statement redefines the file name STATSHR. It becomes the logical name of the shareable image SYS\$SYSDEVICE:[MAIN]STATSHR.EXE for executable images calling STATSHR.

```
$ DEFINE STATSHR SYS$SYSDEVICE:[MAIN]STATSHR
```

Note

Logical names defined in the process or group logical name table are ignored when you run a privileged executable image. Only logical names and table names defined in executive or kernel modes are used to find the image.

17.10.14 Removing Known Images

The INSTALL command REMOVE removes a known file entry for an image and deletes any global sections created when the image was installed. Note that a volume cannot be dismounted while any known file entries are associated with it. To dismount a volume, you must delete all known images associated with it. You must also wait for all processes using those images to exit. Use the DCL command SHOW DEVICES/FILES to determine the status of the files.

For more information about the INSTALL command DELETE, refer to the INSTALL section of the *OpenVMS System Management Utilities Reference Manual*.

17.11 Reserved Memory Registry

The Reserved Memory Registry through its interface within the SYSMAN utility allows an OpenVMS Alpha system to be configured with large amounts of memory set aside for use within memory-resident sections and by other privileged applications. The Reserved Memory Registry also allows an OpenVMS system to be properly tuned through the AUTOGEN utility, taking into account the preallocated reserved memory.

With the Reserved Memory Registry you can:

- Reserve system nonfluid memory for the memory-resident global section fault option.
- Reserve suitable physical pages as well as system non-fluid memory for the memory-resident global section allocate option.

The Reserved Memory Registry includes the ability to specify that the preallocated pages are to be zeroed during the booting of the system. This option reduces the time required to create the memory-resident global demand-zero section.

Another option within the Reserved Memory Registry is to include the size of the page tables required to map to the memory-resident global section in the reserved memory. If this option is specified and the reserved memory is being used for a memory-resident global section, the memory-resident global section is created with shared page tables.

17.11.1 Using the Reserved Memory Registry

OpenVMS provides a mechanism to reserve non-fluid memory for use within a memory-resident global demand-zero section. The reserved memory may either be simply a deduction from the system's nonfluid memory size or pages may be preallocated.

Using the Reserved Memory Registry ensures that AUTOGEN tunes the system properly to not include memory-resident section pages in its calculation of the system's fluid page count. AUTOGEN sizes the system pagefile, number of processes and working set maximum size based on the system's fluid page count. A system can experience severe performance problems if AUTOGEN adjusts parameters based upon a fluid page count that does not account for the physical memory that is permanently reserved for some other purpose.

Using the Reserved Memory Registry also ensures that contiguous, aligned memory is available for memory-resident sections when the allocate option is used.

Note

Although this section describes how to use the reserved memory registry for global sections, this feature can be used for other privileged applications.

17.11.1.1 Reserved Memory Registry Data File

Consumers of reserved, nonfluid memory enter the characteristics of the memory into a data file that is read during the system initialization (boot-time). The mechanics of manipulating the data file are similar to SYS\$LOADABLE_IMAGES:VMS\$SYSTEM_IMAGES.DATA (indicates installation-specific executive loaded images).

This file is called:

```
SYS$SYSTEM:VMS$RESERVED_MEMORY.DATA
```

The file is maintained by the SYSMAN utility (as is the executive loaded image data file).

17.11.1.2 AUTOGEN

The Reserved Memory Registry file, VMS\$RESERVED_MEMORY.DATA, is read by the AUTOGEN feedback mechanism and factors into the setting of the system's fluid page count. AUTOGEN sizes the system pagefile, number of processes and working set maximum size based on the system's fluid page count.

17.11.1.3 Adding Entries to the Reserved Memory Registry

You add an entry to the data file by using the SYSMAN utility. The SYSMAN command is as follows:

```
SYSMAN RESERVED_MEMORY ADD gs_name -  
    /GROUP = n -  
    /SIZE = {size of reserved memory, unit: MB} -  
    /[NO]ALLOCATE -  
    /[NO]ZERO -  
    /[NO]PAGE_TABLES
```

- The gs_name field is the name of the memory-resident global section associated with this reserved memory. A name must be specified.

Performance Considerations

17.11 Reserved Memory Registry

- If the `/GROUP` qualifier is not present, the reserved memory is for a system global section (SYSGBL).
- If the `/GROUP` qualifier is present, the reserved memory is for a group global section. The value *n* UIC group number (in octal) of the process creates the group global section. Only processes within the creator's UIC group number are allowed access to the global section. For example, if a process with the UIC of [6,100] is the creator of the group global section, the group number specified for the `/GROUP` qualifier would be 6.
- If the `/ALLOCATE` qualifier is not specified or if the `/NOALLOCATE` qualifier is specified, the reserved memory is not allocated during the next reboot of the system. The reserved memory is only deducted from the system's fluid page count and the creation of the memory-resident global section results in the fault option being used.
- If the `/ALLOCATE` qualifier is specified, pages are allocated during the next reboot of the system. The allocated memory is deducted from the system's fluid page count and the creation of the memory-resident global section results in the allocate option being used. The physical alignment of the pages is based on the maximum granularity hint factor that can be used to map the pages given the size of the reserved memory. Possible granularity hint factors are 512 pages (or 4 MB) and 64 pages (or 512 KB). Therefore, assuming an 8-KB system page size, reserved memory is physically aligned as follows:
 1. size \geq 4 MB: physically aligned on a 4 Mbyte boundary
 2. size $<$ 4 MB: physically aligned on a 512 KB boundary
- If the `/ZERO` qualifier is not specified or if `/NOZERO` is specified, the preallocated pages are not zeroed during system initialization. The pages are zeroed at the time the global section is created.
- The `/ZERO` qualifier is only allowed if the `/ALLOCATE` qualifier is specified. If the `/ZERO` qualifier is specified, the preallocated pages are zeroed during system initialization. Zeroed pages are required for memory-resident global sections; however, the pages need not be zeroed during system initialization.
- If the `/PAGE_TABLES` qualifier is not specified or if `/NOPAGE_TABLES` is specified, additional memory is not reserved for shared page tables. When the memory-resident global section is created, shared page tables are not created for the global section.
- If the `/PAGE_TABLES` qualifier is specified, additional memory is reserved for shared page tables. When the memory-resident global section is created, shared page tables are created for the global section. If the `/ALLOCATE` qualifier is not specified or if `/NOALLOCATE` is specified, the additional reserved memory is only deducted from the system's fluid page count. If the `/ALLOCATE` qualifier is specified, additional pages are allocated during the next reboot of the system for the shared page tables and the additional reserved memory is deducted from the system's fluid page count.

17.11.2 Removing Entries from the Reserved Memory Registry

You can remove a reserved memory entry by issuing the following SYSMAN command:

```
SYSMAN RESERVED_MEMORY REMOVE gs_name /GROUP = n
```

The specified `gs_name` is the name of the memory-resident section associated with the entry being removed from the Reserved Memory Registry. A name must be specified.

The value `n` specified by the `/GROUP` qualifier is the UIC group number (in octal) associated with the memory-resident section being removed. If the memory-resident global section is a group section, you must specify the `/GROUP` qualifier. If the memory-resident global section is a system global section, you must not specify the `/GROUP` qualifier.

If page tables are reserved for the named memory-resident global section, the additional reserved memory is also removed.

The `REMOVE` command only removes entries from the Reserved Memory Registry data file; it does not affect memory within the running system.

17.11.2.1 Allocating Reserved Memory

During system initialization, the `VMS$RESERVED_MEMORY.DATA` data file is read.

For each entry in the data file, the number of megabytes is deducted from the system's fluid page count for this memory-resident global section as specified by the `/SIZE` qualifier on the `RESERVED_MEMORY ADD` command. If `/PAGE_TABLES` was specified, the amount of memory required for the shared page tables mapping the memory-resident global section is deducted from the system's fluid page count as well.

If `/ALLOCATE` was specified on the `RESERVED_MEMORY ADD` command, a suitable chunk of physical pages is also allocated and set aside for the memory-resident global section. If `/PAGE_TABLES` was specified, an additional chunk of physical pages is allocated and set aside for the shared page tables. The pages have a physical alignment appropriate for the largest granularity hint factor for the given sized chunk. If `/ZERO` was specified, the pages are zeroed during system initialization or when the system is idle. If `/ZERO` was not specified or if `/NOZERO` was specified, the pages are zeroed at the time the memory-resident global section is created.

If the system parameter `STARTUP_P1` is set to `MIN`, entries in the Reserved Memory Registry entries are ignored and memory is not reserved.

If errors are encountered during system initialization while processing the Reserved Memory Registry data file, with reserving system fluid pages, or with allocating contiguous, aligned physical pages, an error message is issued to the console and the system continues to boot.

17.11.2.2 Freeing Reserved Memory

In the running system, you can free reserved memory by issuing the following SYSMAN command:

```
SYSMAN RESERVED_MEMORY FREE gs_name /GROUP = n
```

The specified `gs_name` is the name of the memory-resident section associated with the entry being freed from the Reserved Memory Registry. A name must be specified.

Performance Considerations

17.11 Reserved Memory Registry

The value *n* specified by the /GROUP qualifier is the UIC group number (in octal) associated with the memory-resident section being freed. If the memory-resident global section is a group global section, you must specify the /GROUP qualifier. If the memory-resident global section is a system global section, you must not specify the /GROUP qualifier.

If physical pages were not preallocated during system initialization for this global section, the reserved memory is simply added to the system's fluid page count. Otherwise the physical pages are deallocated onto the system's free or zeroed page list. The system's fluid page count is adjusted to include the deallocated pages.

If page tables are also reserved for the named memory-resident global section, the reserved memory for the shared page tables is also freed.

If the reserved memory is in use by the named memory-resident global section, the amount of reserved memory not currently in use is freed.

The RESERVED_MEMORY FREE command does not affect the Reserved Memory Registry data file contents, it only affects the memory within the running system.

17.11.2.3 Displaying Reserved Memory

Two different places hold reserved memory information, the Reserved Memory Registry data file and the Reserved Memory Registry in the running system created during system initialization based on the entries in the data file.

Different display mechanisms may be used depending on where the information about the reserved memory originates.

There are three mechanisms for displaying the Reserved Memory Registry within the running system: SYSMAN, the DCL SHOW MEMORY command and SDA.

- SYSMAN

You can display the Reserved Memory Registry within the running system by issuing the following SYSMAN command:

```
SYSMAN RESERVED_MEMORY SHOW gs_name /GROUP = n
```

The specified *gs_name* is the name of the memory-resident global section associated with the entry being displayed from within the running system. If *gs_name* is not specified, the reserved memory for all registered global sections is displayed.

The value specified by the /GROUP qualifier is the UIC The value specified by the /GROUP qualifier is the UIC *n* group number (in octal) associated with the memory-resident section being displayed. If the memory-resident global section is a group global section, you must specify the /GROUP qualifier. If the memory-resident global section is a system global section, you must not specify the /GROUP qualifier. The /GROUP qualifier is allowed only if *gs_name* is specified.

- DCL SHOW MEMORY command

You can display the Reserved Memory Registry within the running system by issuing the DCL SHOW MEMORY command. This command shows all memory-related information about the running system including information about the Reserved Memory Registry.

SHOW MEMORY /RESERVED displays just information about the Reserved Memory Registry within the running system.

Performance Considerations

17.11 Reserved Memory Registry

The information displayed by SHOW MEMORY includes how much memory is currently in use by the named global section. It also includes how much memory is reserved and in use (if any) for page tables.

- SDA

SDA also includes various enhancements to display the Reserved Memory Registry within the running system as well as in crash dump files. The command interface is TBD.

17.11.2.4 Using Reserved Memory

The system services SYS\$CREATE_GDZRO and SYS\$CRMPSC_GDZRO_64 call internal kernel mode OpenVMS Alpha routines to use the reserved memory registered in the Reserved Memory Registry.

The global section need not be registered in the Reserved Memory Registry. If the global section name is registered in the Reserved Memory Registry, the size of the global section need not exactly match the size of the reserved memory. If the global section is not registered, or if /NOALLOCATE was specified when the global section was registered in the Reserved Memory Registry, the fault option is used for the memory-resident global DZRO section. If the size is greater than the size of the reserved memory, the system service call to create the memory-resident global DZRO section fails if there are not enough additional fluid pages within the system.

If /ALLOCATE was specified when the global section was registered in the Reserved Memory Registry, the allocate option is used for the memory-resident global DZRO section. The size of the global section must be less than or equal to the size of the reserved, preallocated memory or the error SS\$_MRES_PFN SMALL is returned by the system service call.

17.11.2.5 Returning Reserved Memory

When a memory-resident global section is deleted, the physical pages used for that global section are deallocated to the free page list if physical pages were not preallocated for this global section. The system's fluid page count is adjusted only for those pages not reserved in the Reserved Memory Registry for this global section.

When a memory-resident global section is deleted, the physical pages used for that global section are returned to the Reserved Memory Registry if physical pages were preallocated for this global section. The physical pages are not deallocated to the free page list and are still reserved. No adjustment is made to the system's fluid page count.

Reserved memory may only be freed to the running system by using the RESERVED_MEMORY FREE command to the SYSMAN utility.

Note

Permanent global sections are deleted by calling SYS\$DGBLSC and upon the last reference to the global section. Nonpermanent global sections are simply deleted upon last reference to the global section.

Performance Considerations

17.11 Reserved Memory Registry

17.11.3 Application Configuration

The configuration of an OpenVMS Alpha application that uses memory-resident global sections performs the following steps:

1. Execute the `SYSMAN RESERVED_MEMORY ADD` commands that specify the required use of reserved memory.
2. Run `AUTOGEN` with feedback to set the system's fluid page count appropriately and size the system's pagefile, number of processes and working set maximum size appropriately.
3. Reboot the system to allow for the reserved memory to be deducted from the system's fluid page count and for contiguous, aligned pages to be allocated and zeroed as necessary.

Managing File System Data Caches

This chapter describes how to manage the Extended File Cache (XFC) (Alpha only) and the Virtual I/O Cache (VIOC). These are the caches that the Files-11 file system uses to cache data for ODS-2 and ODS-5 volumes.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Disabling caching clusterwide	Section 18.3
Mounting a volume with caching disabled	Section 18.4
Managing the Extended File Cache (Alpha Only)	Section 18.5
Managing the Virtual I/O Cache	Section 18.6

This chapter explains the following concepts:

Concept	Section
Caching	Section 18.1
File system data caches	Section 18.2

18.1 Understanding Caching

The Files-11 file system uses a technique called **caching** to improve performance. It keeps a copy of data that it has recently read from disk or written to disk in an area of memory called a **cache**.

When an application reads data, the file system checks whether the data is in its cache. It issues an I/O to read the data from disk only if the data is not in the cache.

Caching improves read performance. Reading data from memory (from the cache) is much faster than reading it from disk.

There are several levels of caching in both the hardware I/O subsystem and in OpenVMS. In general, more levels of caching result in better response time in accessing data.

Managing File System Data Caches

18.2 Understanding File System Data Caches

18.2 Understanding File System Data Caches

For ODS-2 and ODS-5 volumes, the Files-11 file system has several caches. It has **metadata caches** for file metadata such as file headers, and a **data cache** for file data. It can use one of two system-wide data caches, as follows:

File System Data Cache	Description
Virtual I/O Cache (VIOC)	This is the original data cache, and is available on VAX and Alpha.
Extended File Cache (XFC)	This data cache is available on OpenVMS Alpha Version 7.3 and later, and is available only on OpenVMS Alpha. It provides better performance and more capability than VIOC. XFC is the default cache on OpenVMS Alpha Version 7.3 and later.

This chapter describes how to manage the data caches. For information on managing the metadata caches, see the *OpenVMS Performance Management*. Note that RMS makes use of local and global buffers that can perform data caching. By default, this caching is not enabled. You can manipulate the RMS local and global buffers to affect I/O performance. For information on managing the RMS local and global buffers, see the *Guide to OpenVMS File Applications*. Note also that the modified page list is a type of cache; that is, if a process references a page that is on the modified page list, the page is placed back into the working set of the process and is not output to disk.

Both XFC and VIOC are virtual block caches that maintain consistent data within an OpenVMS Cluster. They cache both data files and image files. The data caches are write-through caches; when an application writes data to a file, the data is written straight through to disk. The application has to wait until the disk I/O completes and the data is on disk.

In an OpenVMS Cluster, different nodes can use different data caches. This allows mixed architecture clusters to benefit from XFC. OpenVMS Alpha nodes can use either XFC or VIOC. OpenVMS VAX nodes can use only VIOC, as described in Section 18.5.6.

XFC improves I/O performance and contains the following features that are not available with VIOC:

- Read-ahead caching
- Automatic resizing of the cache
- Larger maximum cache size
- No limit on the number of closed files that can be cached
- Control over the maximum size of I/O that can be cached
- Control over whether cache memory is static or dynamic

18.3 Disabling Caching Clusterwide

At system startup, the value of a static system parameter controls whether the file system uses a data cache, and if so, which data cache it uses (XFC or VIOC). The system parameter depends on the operating system, as shown in the following table:

Operating System	System Parameter	Enabled	Disabled
OpenVMS Alpha	VCC_FLAGS	1 or 2 (default) †	0
OpenVMS VAX	VBN_CACHE_S	1 (default)	0

†1 selects VIOC, 2 (the default) selects XFC

In an OpenVMS Cluster, if file system data caching is disabled on one node, it is disabled throughout the cluster. The other nodes in the cluster cannot use XFC or VIOC until that node leaves the cluster or reboots with VCC_FLAGS or VBN_CACHE_S set to a non-zero value. You can use the DCL command SHOW MEMORY to determine whether caching is enabled.

To disable caching clusterwide, follow these steps on any node in your OpenVMS Cluster:

1. Set the appropriate system parameter (VCC_FLAGS or VBN_CACHE_S) to 0, using MODPARAMS.DAT.
2. Run AUTOGEN to ensure that other system parameters allow for the new value. This is not essential, but it is advisable.
3. Reboot the system to make the new value effective.

18.4 Mounting a Volume With Caching Disabled

To prevent the file system from caching data on a particular volume, such as a database volume, use the MOUNT /NOCACHE command to mount the volume with caching disabled.

If you are using XFC in an OpenVMS Cluster, mounting a volume /NOCACHE is easier than using SET FILE /CACHING_ATTRIBUTE to set the caching attribute of all the files in the volume to no caching (see Section 18.5.3). Using MOUNT /NOCACHE ensures the minimum caching overhead. Note that the MOUNT/NOCACHE command also disables XQP caching.

Example

This example mounts a database volume labeled ORACLE_VOL1 with caching disabled:

```
$ MOUNT DUA100: ORACLE_VOL1 /NOCACHE /SYSTEM
```

18.5 Managing XFC (Alpha Only)

This section describes how to manage XFC, which is available only on OpenVMS Alpha. It describes the following tasks.

Managing File System Data Caches

18.5 Managing XFC (Alpha Only)

Task	Section
Controlling the size of the cache	Section 18.5.1
Controlling the maximum cached I/O size	Section 18.5.2
Disabling caching for a file	Section 18.5.3
Disabling read-ahead caching	Section 18.5.4
Monitoring performance	Section 18.5.5
Using in a mixed architecture OpenVMS Cluster	Section 18.5.6

18.5.1 Controlling the Size of the Cache

This section describes how to control the minimum and maximum size of XFC

XFC is held in virtual memory in S2 space and automatically shrinks and grows, depending on your I/O workload and how much spare memory is available on your system. S2 space is a 64-bit address space, so the cache can grow to very large sizes when required.

As your I/O workload increases, the cache automatically grows, but never to more than the maximum size. And when your applications need memory, the cache automatically shrinks, but never to less than the minimum size.

18.5.1.1 Controlling the Minimum Cache Size

The minimum size of XFC is controlled by the value of the `VCC$MIN_CACHE_SIZE` entry in the reserved memory registry. `VCC$MIN_CACHE_SIZE` specifies the amount of memory in megabytes (MB) that is allocated to XFC at system startup. The cache never shrinks below this size. This memory is never released.

Checking the Minimum Size

To check the minimum size of XFC, use either the Sysman utility command `RESERVED_MEMORY /SHOW` or the DCL command `SHOW MEMORY /RESERVED`. For example:

```
$ SHOW MEMORY /RESERVED
System Memory Resources on 11-MAY-2000 15:50:25.64

Memory Reservations (pages):      Group   Reserved   In Use   Type
VCC$MIN_CACHE_SIZE                ---      1536      1536    Allocated
Total (400.00 Mb reserved)        1536      1536
```

Setting a Minimum Size

By default, the reserved memory registry does not contain an entry for `VCC$MIN_CACHE_SIZE`, so no memory is allocated to XFC at system startup. However, XFC allocates a small amount of memory to maintain overall system throughput. The amount of memory allocated varies according to the size of your machine.

To set a minimum size, follow these steps:

1. Use the Sysman utility's `RESERVED_MEMORY ADD` command to add an entry for `VCC$MIN_CACHE_SIZE`. For example, to set the minimum size to 300MB:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> RESERVED MEMORY ADD VCC$MIN CACHE SIZE /SIZE=300 /ALLOCATE -
SYSMAN> /NOGLOBAL_SECTION /NOZERO /NOPAGE_TABLE
```

You must use all the qualifiers shown in this example. For best performance, set the minimum size to a multiple of 4MB.

Managing File System Data Caches

18.5 Managing XFC (Alpha Only)

Note that if you are doing this on a NUMA type machine, where you want to allocate reserved memory in different RADs, you may choose to use the following commands (rather than those above) to set the minimum size:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> RESERVED MEMORY ADD VCC$MIN CACHE SIZE /SIZE=300 /ALLOCATE -
SYSMAN> /NOGLOBAL SECTION /NOZERO /NOPAGE TABLE/RAD=0
SYSMAN> RESERVED MEMORY EXTEND VCC$MIN CACHE SIZE /SIZE=500 /ALLOCATE -
SYSMAN> /NOGLOBAL SECTION /NOZERO /NOPAGE TABLE/RAD=1
```

2. Run AUTOGEN to ensure that other system parameters allow for the new value. This is not essential, but it is advisable.
3. Reboot the system to make the new value effective.

During startup, if the system does not have enough memory to allocate the specified minimum size, no memory is allocated to XFC and its minimum size is set to 0MB.

Changing the Minimum Size

To change the minimum size of XFC, follow these steps:

1. Use the Sysman utility's RESERVED_MEMORY MODIFY command to modify the existing entry for VCC\$MIN_CACHE_SIZE. For example, to change the minimum size to 360MB:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> RESERVED MEMORY MODIFY VCC$MIN_CACHE_SIZE /SIZE=360 /ALLOCATE -
SYSMAN> /NOGLOBAL SECTION /NOZERO
```

You must use all the qualifiers shown in this example. For best performance, remember to set the minimum size to a multiple of 4MB.

2. Run AUTOGEN to ensure that other system parameters allow for the new value. This is not essential, but it is advisable.
3. Reboot the system to make the new value effective.

During startup, if the system does not have enough memory to allocate the specified minimum size, no memory is allocated to XFC and its minimum size is set to 0MB.

18.5.1.2 Controlling the Maximum Cache Size

To control the maximum size of XFC, use the dynamic system parameter VCC_MAX_CACHE. It specifies the size in megabytes.

By default, VCC_MAX_CACHE is -1, which means that at system startup, the maximum size of XFC is set to 50 percent of the physical memory on the system. For example, if the system has 2GB of physical memory, its maximum size is set to 1GB.

Note that the maximum size specified by VCC_MAX_CACHE does not include the memory that XFC consumes indirectly via the OpenVMS lock manager.

The value of VCC_MAX_CACHE at system startup sets an upper limit for the maximum size of XFC. You cannot increase the maximum size beyond that value.

For example, VCC_MAX_CACHE is 60 at system startup, so the maximum size is initially set to 60MB. You then set VCC_MAX_CACHE to 40, which decreases the maximum size to 40MB. If XFC is bigger than 40MB, it gradually shrinks to 40MB. You then set VCC_MAX_CACHE to 80, but the maximum size is only increased to 60MB, the value set at system startup. You cannot increase the maximum size beyond the value set at system startup.

Managing File System Data Caches

18.5 Managing XFC (Alpha Only)

If `VCC_MAX_CACHE` is less than the minimum size specified by the value of the `VCC$MIN_CACHE_SIZE` entry in the reserved memory registry, then during system startup, `VCC_MAX_CACHE` is ignored and the maximum size of XFC is set to the same value as the minimum size. In this case, XFC has a fixed size and cannot shrink or grow.

Example

This example changes the maximum size of XFC to 50MB:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> USE CURRENT
SYSGEN> SET VCC MAX CACHE 50
SYSGEN> WRITE ACTIVE
```

18.5.1.3 Enabling a Static Cache Size

On larger machines, XFC may be somewhat limited by the default size of the modified page list. In general, the modified page list can be considered a 1 to 1 correspondence data cache. XFC is a many to one cache; that is, in general, a cached page is accessed by many users. On large memory systems, AUTOGEN usually sets `MPW_HILIMIT` to a very large value. This may mean that there are not enough free pages for the memory management subsystem to give to XFC.

You can force XFC to have a minimum number of pages as specified in Setting a Minimum Size. You can also permanently allocate memory just for XFC, which prevents any overhead caused by the dynamic allocation and deallocation routines (similar in operation to `VIOC`). To do this, set system parameter `VCC_MAX_CACHE` to be equal to the memory reservation specified by `VCC$MIN_CACHE_SIZE`.

For example, if your system has 128GB of memory, you may find that by dedicating 8GB of memory to XFC, your cache hit rates and overall response time is consistently good. For example:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> USE CURRENT
SYSGEN> SET VCC MAX CACHE 8000
SYSGEN> WRITE ACTIVE
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> RESERVED MEMORY ADD VCC$MIN CACHE SIZE /SIZE=8000 /ALLOCATE -
  _SYSMAN> /NOGLOBAL_SECTION /NOZERO /NOPAGE_TABLE/RAD=0
```

18.5.2 Controlling the Maximum Cached I/O Size

The dynamic system parameter `VCC_MAX_IO_SIZE` controls the maximum size of I/O that can be cached by XFC. This parameter specifies the size in blocks. By default, it is 128.

Example

This example changes the maximum size of I/O that can be cached by XFC to 1,000 blocks:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> USE CURRENT
SYSGEN> SET VCC MAX IO_SIZE 1000
SYSGEN> WRITE ACTIVE
```

This series of commands affects I/Os to volumes currently mounted on the local node, as well as to volumes mounted in the future. After you enter these commands, XFC does not cache I/Os that are bigger than 1,000 blocks. The `SHOW MEMORY /CACHE /FULL` command provides a histogram of I/O sizes that can provide guidelines for efficient parameter setting.

18.5.3 Disabling Caching for a File

To prevent XFC from caching a particular file, such as a database file, set the caching attribute of the file to no caching.

The **caching attribute** of a file is the default caching option that is used by XFC when an application accesses the file without specifying which caching option it wants to use. The caching option can be either write-through caching or no caching.

If you want a file to be cached, set its caching attribute to write-through (the default). If you don't want a file to be cached, set its caching attribute to no caching.

Use...	To...
SET FILE /CACHING_ATTRIBUTE=keyword †	Set the caching attribute of a file or directory
DIRECTORY /CACHING_ATTRIBUTE or DIRECTORY /FULL	Show the caching attribute of a file or directory
DIRECTORY /SELECT=CACHING_ ATTRIBUTE=(keyword[,...]) †	Show all the files and directories that have a particular caching attribute

†keyword can be either WRITETHROUGH or NO_CACHING.

XFC does not cache directories. The caching attribute of a directory controls only how the caching attribute is inherited by new files and subdirectories created in the directory:

- When you create a new directory or file, it inherits its caching attribute from its parent directory.
- When you create a new version of an existing file, the new file inherits its caching attribute from the highest version of the existing file.

Examples

- This example sets the caching attribute to no caching for all the files in and under the directory [SMITH.BORING].

The first SET FILE command sets the attribute for the directory to make sure that all files and subdirectories subsequently created in it inherit the attribute. The second SET FILE command sets the attribute for all existing files and directories in and under the directory.

```
$ SET FILE DISK$USERS:[SMITH]BORING.DIR;1 /CACHING_ATTRIBUTE=NO_CACHING
$ SET FILE DISK$USERS:[SMITH.BORING...] *.* /CACHING_ATTRIBUTE=NO_CACHING
```

- This example uses the DIRECTORY command's /CACHING_ATTRIBUTE qualifier to show the caching attribute of MYFILE.TXT:

```
$ DIRECTORY MYFILE.TXT /CACHING_ATTRIBUTE
Directory DISK$USERS:[SMITH]
MYFILE.TXT;1          Write-through
Total of 1 file.
```

- This example shows all the files in the volume DISK\$USERS that have a caching attribute of no caching.

```
$ DIRECTORY DISK$USERS:[000000...] *.* /SELECT=CACHING_ATTRIBUTE=NO_CACHING
```

Managing File System Data Caches

18.5 Managing XFC (Alpha Only)

18.5.4 Disabling Read-Ahead Caching

XFC uses a technique called **read-ahead caching** to improve the performance of applications that read data sequentially. It detects when a file is being read sequentially in equal-sized I/Os, and fetches data ahead of the current read, so that the next read instruction can be satisfied from cache.

To disable read-ahead caching on the local node, set the dynamic system parameter `VCC_READAHEAD` to 0. By default, this parameter is 1, which allows the local node to use read-ahead caching.

Example

This example disables read-ahead caching on the local node:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> USE CURRENT
SYSGEN> SET VCC_READAHEAD 0
SYSGEN> WRITE ACTIVE
```

This series of commands affects volumes currently mounted on the local node, as well as volumes mounted in the future. Once you enter these commands, read-ahead caching is not used on the local node.

18.5.5 Monitoring Performance

XFC provides more information than VIOC. For example, you can obtain information on system-wide, volume-wide, or even a per-file basis. Disk I/O response times are also available. See the *OpenVMS DCL Dictionary: N-Z* for a description of the `SHOW MEMORY` command.

18.5.5.1 System-Wide Statistics

Use `SHOW MEMORY /CACHE` to monitor the overall system performance of XFC. For example:

```
$ SHOW MEMORY /CACHE

System Memory Resources on 26-JAN-2001 15:58:18.71

Extended File Cache (Time of last reset: 24-JAN-2001 15:03:39.05)
Allocated (Mbytes)      ①      3000.00  Maximum size (Mbytes) ⑪      5120.00
Free (Mbytes)          ②      2912.30  Minimum size (Mbytes) ⑫      3000.00
In use (Mbytes)        ③      87.69   Percentage Read I/Os ⑬      98%
Read hit rate          ④      92%    Write hit rate        ⑭      0%
Read I/O count         ⑤      178136  Write I/O count       ⑮      1867
Read hit count         ⑥      165470  Write hit count       ⑯      0
Reads bypassing cache ⑦      2802   Writes bypassing cache ⑰      39
Files cached open     ⑧      392    Files cached closed  ⑱      384
Vols in Full XFC mode ⑨      0      Vols in VIOC Compatible mode ⑲      4
Vols in No Caching mode ⑩      1      Vols in Perm. No Caching mode ⑳      0
```

- | | |
|-----------------|--|
| ① Allocated | The amount of memory currently allocated to the cache. |
| ② Free | The amount of memory currently allocated to the cache that is not being used. |
| ③ In Use | The amount of memory currently allocated to the cache that is being used. This is the difference between the Allocated value and the Free value. |
| ④ Read hit rate | The ratio of the Read hit count field divided by the Read I/O count field. |

Managing File System Data Caches

18.5 Managing XFC (Alpha Only)

5	Read I/O count	The total number of read I/Os seen by the cache since system startup.
6	Read hit count	The total number of read hits since system startup. A read hit is a read I/O that did not require a physical I/O to disk because the data was found in the cache.
7	Reads bypassing cache	The total number of read I/Os since system startup that were seen by the cache but were not cached, for example, because they were too big, or they were for volumes mounted /NOCACHE, or they specified one of the following QIO modifiers: IO\$M_DATACHECK, IO\$M_INHRETRY, or IO\$M_NOVCACHE.
8	Files cached open	The number of open files currently being cached.
9	Volumes in Full XFC mode	Reserved for Compaq. Should be 0.
10	Volumes in No Caching mode	If caching is disabled on the local node or on another node in the OpenVMS Cluster, this is the number of volumes that are currently mounted on the local node. Otherwise it is zero.
11	Maximum size	The maximum size that the cache could ever grow to.
12	Minimum size	The minimum size that the cache could ever shrink to. This is controlled by the value of the VCC\$MIN_CACHE_SIZE entry in the reserved memory registry.
13	Percentage Read I/Os	Percentage of I/Os that are reads.
14	Write hit rate	This field is reserved for future use.
15	Write I/O count	The total number of write I/Os seen by the cache since system startup.
16	Write hit count	This field is reserved for future use.
17	Writes bypassing cache	The total number of write I/Os since system startup that were seen by the cache but were not cached, for example, because they were too big, or they were for volumes mounted /NOCACHE, or they specified one of the following QIO modifiers: IO\$M_DATACHECK, IO\$M_ERASE, IO\$M_INHRETRY, or IO\$M_NOVCACHE.
18	Files cached closed	The number of closed files that still have valid data in the cache.
19	Volumes in VIOC compatible mode	The number of volumes being cached by XFC that are using the VCC caching protocol. As of OpenVMS Version 7.3, XFC uses only VCC caching protocol.
20	Vols in Perm. No Caching mode	This field should be zero. If nonzero, XFC has detected an illegal write operation to this device and has disabled caching to this device.

See the *OpenVMS DCL Dictionary: N-Z* for a description of the SHOW MEMORY command.

18.5.6 Using XFC in a Mixed Architecture OpenVMS Cluster

In an OpenVMS Cluster, some nodes can use XFC and other nodes can use VIOC. This allows mixed architecture clusters to benefit from XFC.

When a volume is mounted on a node that is using VIOC, the nodes using XFC cannot cache any files in the volume that are shared for writing. A file that is **shared for writing** is one that is being accessed by more than one node in an OpenVMS Cluster, and at least one of those nodes opened it for write access.

Managing File System Data Caches

18.6 Managing the Virtual I/O Cache

18.6 Managing the Virtual I/O Cache

This section describes how to manage VIOC. It describes the following tasks:

Task	Section
Selecting VIOC on an Alpha system	Section 18.6.2
Controlling the size of the cache	Section 18.6.3
Monitoring performance	Section 18.6.4

The virtual I/O cache is a clusterwide, write-through, file-oriented, disk cache that can reduce the number of disk I/O operations and increase performance. The purpose of the virtual I/O cache is to increase system throughput by reducing file I/O response times with minimum overhead. The virtual I/O cache operates transparently of system management and application software, and maintains system reliability while it significantly improves virtual disk I/O read performance.

18.6.1 Understanding How the Cache Works

The virtual I/O cache can store data files and image files. For example, ODS-2 disk file data blocks are copied to the virtual I/O cache the first time they are accessed. Any subsequent read requests of the same data blocks are satisfied from the virtual I/O cache (hits) eliminating any physical disk I/O operations (misses) that would have occurred.

Depending on your system work load, you should see increased application throughput, increased interactive responsiveness, and reduced I/O load.

Note

Applications that initiate single read and write requests do not benefit from virtual I/O caching as the data is never reread from the cache. Applications that rely on implicit I/O delays might abort or yield unpredictable results.

Several policies govern how the cache manipulates data, as follows:

- Write-through—All write I/O requests are written to the cache as well as to the disk.
- Least Recently Used (LRU)—If the cache is full, the least recently used data in the cache is replaced.
- Cached data maintained across file close—Data remains in the cache after a file is closed.
- Allocate on read and write requests—Cache blocks are allocated for read and write requests.

18.6.2 Selecting VIOC on an Alpha System

If for some reason, you want an Alpha system to use VIOC instead of XFC, follow these steps:

1. Remove the entry for VCC\$MIN_CACHE_SIZE from the reserved memory registry, using the Sysman utility's RESERVED_MEMORY REMOVE command:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> RESERVED_MEMORY REMOVE VCC$MIN_CACHE_SIZE /NOGLOBAL_SECTION
```

This makes sure that no memory is allocated to XFC in Step 4, when the system reboots with VIOC.

2. Set the VCC_FLAGS system parameter to 1.
3. Run AUTOGEN to ensure that other system parameters allow for the new value. This is not essential, but it is advisable.
4. Reboot the system. VIOC is automatically loaded during startup instead of XFC, because VCC_FLAGS is 1.

If you forgot to remove the VCC\$MIN_CACHE_SIZE entry from the reserved memory registry in Step 1, memory is allocated to XFC even though XFC is not loaded. Nothing can use this memory. If this happens, use the Sysman utility's RESERVED_MEMORY FREE command to release this memory:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> RESERVED_MEMORY FREE VCC$MIN_CACHE_SIZE /NOGLOBAL_SECTION
```

18.6.3 Controlling the Size of the Cache

The way that you control the size of VIOC depends on whether you have an OpenVMS Alpha or OpenVMS VAX system.

OpenVMS Alpha

On OpenVMS Alpha systems, the size of VIOC is fixed at system startup time. The cache can't shrink or grow. The value of the static system parameter VCC_MAXSIZE specifies the size of the cache in blocks. By default it is 6400 blocks (3.2MB).

To change the size of VIOC on an OpenVMS Alpha system, follow these steps:

1. Set the VCC_MAXSIZE system parameter to the required value.
2. Run AUTOGEN to ensure that other system parameters allow for the new value. This is not essential, but it is advisable.
3. Reboot the system to make the new value effective.

OpenVMS VAX

On OpenVMS VAX systems, you can use the static system parameter VCC_PTES to specify the maximum size of VIOC. This parameter specifies the size in pages. By default it is 2,000,000,000.

VIOC automatically shrinks and grows, depending on your I/O workload and how much spare memory is available on your system. As your I/O workload increases, the cache automatically grows, but never to more than the maximum size. And when your applications need memory, the cache automatically shrinks.

Managing File System Data Caches

18.6 Managing the Virtual I/O Cache

To change the maximum size of VIOC on an OpenVMS VAX system, follow these steps:

1. Set the VCC_MAXSIZE system parameter to the required value.
2. Run AUTOGEN to ensure that other system parameters allow for the new value. This is not essential, but it is advisable.
3. Reboot the system to make the new value effective.

18.6.4 Displaying VIOC Statistics

Use the DCL command `SHOW MEMORY/CACHE/FULL` to display statistics about the virtual I/O cache, as shown in the following example:

```
$ SHOW MEMORY/CACHE/FULL

                System Memory Resources on 10-OCT-1994 18:36:12.79
Virtual I/O Cache
Total Size (pages)  ①  2422   Read IO Count           ⑥  9577
Free Pages          ②   18    Read Hit Count         ⑦  5651
Pages in Use       ③  2404   Read Hit Rate          ⑧  59%
Maximum Size (SPTes) ④ 11432  Write IO Count         ⑨  2743
Files Retained     ⑤   99    IO Bypassing the Cache ⑩  88
```

Note

This example shows the output for the `SHOW MEMORY/CACHE/FULL` command on a VAX system. The `SHOW MEMORY/CACHE/FULL` command displays slightly different fields on an Alpha system.

- | | |
|-------------------|--|
| ① Total Size | Displays the total number of system memory pages that VIOC currently controls. |
| ② Free Pages | Displays the number of pages controlled by VIOC that do not contain cache data. |
| ③ Pages in Use | Displays the number of pages controlled by VIOC that contain valid cached data. |
| ④ Maximum Size | Shows the maximum size that the cache could ever grow to. |
| ⑤ Files Retained | Displays the number of files that are closed but the file system control information is being retained because they have valid data residing in the cache. |
| ⑥ Read I/O Count | Displays the total number of read I/Os that have been seen by VIOC since the last system. |
| ⑦ Read Hit Count | Displays the total number of read I/Os that did not do a physical I/O because the data for them was found in the cache since the last system BOOT. |
| ⑧ Read Hit Rate | Displays the read hit count and read I/O count ratio. |
| ⑨ Write I/O Count | Shows the total number of write I/Os that have been seen by the cache since the last system BOOT. |
| ⑩ I/O Bypassing | Displays the count of I/Os that for some reason did not attempt to satisfy the request/update by the cache. |

18.6.5 Enabling VIOC

By default, virtual I/O caching is enabled. Use the following system parameters to enable or disable caching. Change the value of the parameters in MODPARAMS.DAT, as follows:

Parameter	Enabled	Disabled
VCC_FLAGS (Alpha)	1	0
VBN_CACHE_S (VAX)	1	0

Once you have updated MODPARAMS.DAT to change the value of the appropriate parameter, you must run AUTOGEN and reboot the node or nodes on which you have enabled or disabled caching. Caching is automatically enabled or disabled during system initialization. No further user action is required.

18.6.6 Determining If VIOC Is Enabled

SHOW MEMORY/CACHE indicates whether VIOC caching is on or off on a running system. (This is a lot easier than using SYSGEN.)

SYSGEN can be used to examine parameters before a system is booted. For example, you can check the system parameter VCC_FLAGS (on Alpha) or VBN_CACHE_S (on VAX) to see if virtual I/O caching is enabled by using SYSGEN, as shown in the following Alpha example:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> SHOW VCC_FLAGS
```

A value of 0 indicates that caching is disabled; the value 1 indicates caching is enabled.

18.6.7 Memory Allocation and VIOC

The memory allocated to caching is determined by the size of the free-page list. The size of the virtual I/O cache can grow if one of the following conditions is true:

- If the amount of available free memory is twice the value of FREEGOAL and if proactive memory reclamation is enabled for periodically waking processes.
- If the amount of available free memory is equal to the value of FREEGOAL and if proactive memory reclamation is enabled for long-waiting processes.
- If the amount of available free memory is greater than GROWLIM and if proactive memory reclamation is not enabled.

The cache size is also limited by the following:

- The number of system page table entries (SPTE) that are available. This number is a calculated value determined at boot time.
- The demands of the memory management subsystem. The memory management subsystem has a direct interface to cache so that, when necessary, it can demand that the cache return space to it.

How is memory reclaimed from the cache? The swapper can reclaim memory allocated to the virtual I/O cache by using first-level trimming. In addition, a heuristic primitive shrinks the cache returning memory in small increments.

Managing File System Data Caches

18.6 Managing the Virtual I/O Cache

18.6.8 Adjusting VIOC Size

The size of the virtual I/O cache is controlled by the system parameter `VCC_MAXSIZE`. The amount of memory specified by this parameter is statically allocated at system initialization and remains owned by the virtual I/O cache.

To increase or decrease the size of the cache, modify `VCC_MAXSIZE` and reboot the system.

18.6.9 VIOC and OpenVMS Cluster Configurations

The cache works on all supported configurations from single-node systems to large mixed-interconnect OpenVMS Cluster systems. The virtual I/O cache is nodal; that is, the cache is local to each OpenVMS Cluster member. Any base system can support virtual I/O caching; an OpenVMS Cluster license is not required to use the caching feature.

Note

If any member of an OpenVMS Cluster does not have caching enabled, then no caching can occur on any node in the OpenVMS Cluster (including the nodes that have caching enabled). This condition remains in effect until the node or nodes that have caching disabled either enable caching or leave the cluster.

The lock manager controls cache coherency. The cache is flushed when a node leaves the OpenVMS Cluster. Files opened on two or more nodes with write access on one or more nodes are not cached.

Testing the System with UETP

This chapter explains how to use UETP (user environment test package) to test whether the OpenVMS operating system is installed correctly.

19.1 Overview

This overview summarizes what UETP does and how you use it. The rest of the chapter provides detailed instructions for setting up your system for testing, running the tests, and troubleshooting errors.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Running UETP (a summary)	Section 19.1.2
Preparing to use UETP	Section 19.2
Setting up the devices to be tested	Section 19.3
Starting UETP	Section 19.4
Stopping a UETP operation	Section 19.5
Troubleshooting: identifying and solving problems	Section 19.7

This chapter explains the following concepts:

Concept	Section
Understanding UETP	Section 19.1.1
Troubleshooting (an overview)	Section 19.6
UETP Tests and Phases	Section 19.8

19.1.1 Understanding UETP

UETP is a software package designed to test whether the OpenVMS operating system is installed correctly. UETP puts the system through a series of tests that simulate a typical user environment by making demands on the system that are similar to demands that can occur in everyday use.

UETP is not a diagnostic program; it does not attempt to test every feature exhaustively. When UETP runs to completion without encountering nonrecoverable errors, the system being tested is ready for use.

UETP exercises devices and functions that are common to all OpenVMS systems, with the exception of optional features such as high-level language compilers. The system components tested include the following ones:

Testing the System with UETP

19.1 Overview

- Most standard peripheral devices
- System's multiuser capability
- DECnet for OpenVMS software
- Clusterwide file access and locks

19.1.2 Summary of How to Use UETP

This section summarizes the procedure for running all phases of UETP with default values. If you are familiar with the test package, refer to this section. If you want additional information, refer to Section 19.2.

Note

If you are using UETP on an OpenVMS Alpha system, you must execute the `CREATE_SPECIAL_ACCOUNTS.COM` command procedure to create the `SYSTEST` and `SYSTEST_CLIG` accounts before you begin the following procedure. For complete information about the `CREATE_SPECIAL_ACCOUNTS.COM` command procedure, see Section 7.4.

1. Log in to the `SYSTEST` account as follows:

Username: `SYSTEST`
Password:

Caution

Because the `SYSTEST` and `SYSTEST_CLIG` accounts have privileges, unauthorized use of these accounts can compromise the security of your system.

2. Make sure no user programs are running and no user volumes are mounted.

Caution

By design, UETP assumes and requests the exclusive use of system resources. If you ignore this restriction, UETP can interfere with applications that depend on these resources.

3. After you log in, check all devices to be sure that the following conditions exist:
 - All devices you want to test are powered up and are on line to the system.
 - Scratch disks are mounted and initialized.
 - Disks contain a directory named `[SYSTEST]` with `OWNER_UIC=[1,7]`. (You can create this directory with the DCL command `CREATE/DIRECTORY`.)
 - Scratch magnetic tape reels are *physically* mounted on each drive you want tested and are initialized with the label `UETP` (using the DCL command `INITIALIZE`). Make sure magnetic tape reels contain at least 600 feet of tape.

- Scratch tape cartridges have been inserted in each drive you want to test and are initialized with the label UETP (using the DCL command INITIALIZE).
- Line printers and hardcopy terminals have plenty of paper.
- Terminal characteristics and baud rate are set correctly. (Refer to the user's guide for your terminal.)

Note that some communication devices need to be set up by a Compaq support representative. (See Section 19.3.)

If you encounter any problems in preparing to run UETP, read Section 19.3 before proceeding.

4. To start UETP, enter the following command and press Return:

```
$ @UETP
```

UETP responds with the following question:

```
Run "ALL" UETP phases or a "SUBSET" [ALL]?
```

Press Return to choose the default response enclosed in brackets. UETP responds with the following sequence of questions:

```
How many passes of UETP do you wish to run [1]?  
How many simulated user loads do you want [4]?  
Do you want Long or Short report format [Long]?
```

Press Return after each prompt. After you answer the last question, UETP initiates its entire sequence of tests, which run to completion without further input. The final message should look like the following one:

```
*****  
*                                     *  
*      END OF UETP PASS 1 AT 22-JUN-1998 16:30:09.38      *  
*                                     *  
*****
```

Note

If you want to run UETP without using the default responses, refer to Section 19.4, which explains your options.

5. After UETP runs, check the log files for errors. If testing completes successfully, the OpenVMS operating system is in proper working order. If UETP does not complete successfully, refer to Section 19.6 for information about troubleshooting.

Note

After a run of UETP, you should run the Error Log utility to check for hardware problems that can occur during a run of UETP. For information about running the Error Log utility, refer to the *OpenVMS System Management Utilities Reference Manual*.

Testing the System with UETP

19.2 Preparing to Use UETP

19.2 Preparing to Use UETP

This section contains detailed instructions for running UETP, including:

- Logging in
- Using the [SYSTEST] directory

19.2.1 Logging In

Obtain the SYSTEST password from your system manager. Log in to the SYSTEST account from the console terminal as follows:

```
Username: SYSTEST
Password:
```

Note

Because SYSTEST has privileges, unauthorized use of this account can compromise the security of your system.

UETP will fail if you do not run the test from the SYSTEST account. Also, if you try to run UETP from a terminal other than the console terminal, the device test phase displays an error message stating that the terminal you are using is unavailable for testing. You can ignore this message.

After you log in to the SYSTEST account, enter the command SHOW USERS to make sure no user programs are running and no user volumes are mounted. UETP requires exclusive use of system resources. If you ignore this restriction, UETP can interfere with applications that depend on these resources.

Note

The information contained in Section 19.7.2 can help you identify and solve problems, including wrong quotas, privileges, or accounts, that could occur when you are running UETP. Refer to this section before you run UETP.

19.2.2 Using the SYSTEST Directories

If you logged in successfully, your default directory is [SYSTEST] on the system disk. UETP uses this directory to hold all the files used by UETP command procedure (UETP.COM) and temporary files used by UETP during testing.

On a typical system, the DCL command SHOW LOGICAL displays the translation of the logical name SYS\$TEST:

```
$ SHOW LOGICAL SYS$TEST
  "SYS$TEST" = "SYS$SYSROOT:[SYSTEST]" (LNM$SYSTEM_TABLE)
```

To use UETP to test a particular disk, such as a scratch disk, create either a [SYSTEST] directory or a [SYS0.SYSTEST] directory on that disk. Section 19.3.3 discusses setting up scratch disks for testing.

19.3 Setting Up the Devices to Be Tested

After you log in, set up the devices on the system for UETP testing, as described in the following sections. Note that your system might not have all the devices described in this section.

19.3.1 Check Your Devices

Examine all devices that UETP will use to be sure that the following conditions exist:

- All devices you want to test are turned on and are on line.
- Scratch disks are initialized and mounted.
- Disks contain a directory named [SYSTEST] with OWNER_UIC=[1,7]. Use the CREATE/DIRECTORY command if the [SYSTEST] directory does not exist on the disk.
- Scratch magnetic tape reels are *physically* mounted on each drive you want tested and are initialized with the label UETP (using the DCL command INITIALIZE). Make sure magnetic tape reels contain at least 600 feet of tape.
- Scratch tape cartridges have been inserted in each drive you want to test and are mounted and initialized with the label UETP (using the DCL command INITIALIZE).
- Line printers and hardcopy terminals have plenty of paper.
- Terminal characteristics and baud rate are set correctly. (Refer to the user's guide for your terminal).

Note that some communications devices discussed in this section must be set up by a Compaq support representative.

19.3.2 System Disk Space Required

Before running UETP, be sure that the system disk has at least 1200 blocks available. Note that systems running more than 20 load test processes can require a minimum of 2000 available blocks. If you run multiple passes of UETP, log files will accumulate in the default directory and further reduce the amount of disk space available for subsequent passes.

If disk quotas are enabled on the system disk, disable them before you run UETP.

19.3.3 How UETP Works on Disks

The disk test phase of UETP uses most of the available free space on each testable disk in the following manner:

- On each testable disk, the device test phase tries to create two files. The size of these files depends on how much free space is available on the disk. Usually the test creates each file with 0.1% of the free space on the disk. However, if the disk is almost full, the test creates files that are 5 blocks. If the test cannot create 5 block files, it fails. Only the initial file creation can cause the device test to fail because it lacks disk space.
- The test randomly reads and writes blocks of data to the files. After every multiple of 20 writes for each file, the test tries to extend the file. The size of this extension is either 5% of the free disk space or 5 blocks if the file was created with 5 blocks. This process of extension continues until the combined space of the files reaches 75% of the free disk space.

Testing the System with UETP

19.3 Setting Up the Devices to Be Tested

By creating and extending fragmented files in this way, UETP exercises the disk. This allows the test to check for exceeded quotas or a full disk, and to adjust for the amount of available disk space.

As with other disks, shadow sets and volume sets can be tested with UETP; the expectation is that the individual members will be listed as untestable during UETINIDEV (initialization of UETP). UETINIDEV lists errors when testing using a shadow set during the system disk (UETDISK00) pass, however, the shadow set is listed as testable. When testing using a volume set, errors will be noted against all but relative volume number 1, and all but relative volume 1 will be listed as untestable at the end of UETINIDEV.

19.3.4 Prepare Disk Drives

To prepare each disk drive in the system for UETP testing, use the following procedure:

1. Place a scratch disk in the drive and spin up the drive. If a scratch disk is not available, use any disk with a substantial amount of free space; UETP does not overwrite existing files on any volume. If your scratch disk contains files that you want to keep, do not initialize the disk; go to step 3.
2. If the disk does not contain files you want to save, initialize it. For example:

```
$ INITIALIZE DUA1: TEST1
```

This command initializes DUA1 and assigns the volume label TEST1 to the disk. All volumes must have unique labels.

3. Mount the disk. For example:

```
$ MOUNT/SYSTEM DUA1: TEST1
```

This command mounts the volume labeled TEST1 on DUA1. The /SYSTEM qualifier indicates that you are making the volume available to all users on the system.

4. UETP uses the [SYSTEST] directory when testing the disk. If the volume does not contain the directory [SYSTEST], you must create it. For example:

```
$ CREATE/DIRECTORY/OWNER_UIC=[1,7] DUA1:[SYSTEST]
```

This command creates a [SYSTEST] directory on DUA1 and assigns a user identification code (UIC) of [1,7]. The directory must have a UIC of [1,7] to run UETP.

If the disk you have mounted contains a root directory structure, you can create the [SYSTEST] directory in the [SYS0.] tree.

19.3.5 Magnetic Tape Drives

Set up magnetic tape drives that you want to test by performing the following steps:

1. Place a scratch magnetic tape with at least 600 feet of magnetic tape in the tape drive. Make sure that the write-enable ring is in place.
2. Position the magnetic tape at the BOT (beginning-of-tape) and put the drive on line.

Testing the System with UETP

19.3 Setting Up the Devices to Be Tested

3. Initialize each scratch magnetic tape with the label UETP. For example, if you have physically mounted a scratch magnetic tape on MUA1, enter the following command and press Return:

```
$ INITIALIZE MUA1: UETP
```

Magnetic tapes must be labeled UETP to be tested. As a safety feature, UETP does not test tapes that have been mounted with the MOUNT command.

If you encounter a problem initializing the magnetic tape or if the test has a problem accessing the magnetic tape, refer to the description of the INITIALIZE command in the *OpenVMS DCL Dictionary*.

19.3.6 Tape Cartridge Drives

To set up tape cartridge drives you want to test, perform the following steps:

1. Insert a scratch tape cartridge in the tape cartridge drive.
2. Initialize the tape cartridge. For example:

```
$ INITIALIZE MUA0: UETP
```

Tape cartridges must be labeled UETP to be tested. As a safety feature, UETP does not test tape cartridges that have been mounted with the MOUNT command.

If you encounter a problem initializing the tape cartridge, or if the test has a problem accessing the tape cartridge, refer to the description of the DCL INITIALIZE command in the *OpenVMS DCL Dictionary*.

TLZ04 Tape Drives

During the initialization phase, UETP sets a time limit of 6 minutes for a TLZ04 unit to complete the UETTAPE00 test. If the device does not complete the UETTAPE00 test within the allotted time, UETP displays a message similar to the following one:

```
-UETP-E-TEXT, UETTAPE00.EXE testing controller MKA was stopped ($DELPRC)
  at 16:23:23.07 because the time out period (UETP$INIT TIMEOUT)
  expired or because it seemed hung or because UETINIT01 was aborted.
```

To increase the timeout value, enter a command similar to the following one before running UETP:

```
$ DEFINE/GROUP UETP$INIT_TIMEOUT "0000 00:08:00.00"
```

This example defines the initialization timeout value to 8 minutes.

19.3.7 Compact Disc Drives

To run UETP on an RRD40 or RRD50 compact disc drive, you must first load the test disc that you received with your compact disc drive unit.

19.3.8 Optical Disk Drives

To run UETP on an RV60 drive, set up the RV64 optical disk-storage system, perform the following steps:

1. Use the Jukebox Control Software (JCS) to load an optical disk in each of the RV60 drives. JCS is a layered product on the OpenVMS operating system that comes with the RV64 and is responsible for controlling the robot arm that loads and unloads the disks.
2. Initialize the optical disks with the label UETP, but do not mount them.

Testing the System with UETP

19.3 Setting Up the Devices to Be Tested

UETP tests all the RV60s present in the RV64 simultaneously. Unlike the tape tests, UETP does not reinitialize the optical disks at the end of the test.

19.3.9 Terminals and Line Printers

Terminals and line printers must be turned on and on line to be tested by UETP. Check that line printers and hardcopy terminals have enough paper. The amount of paper required depends on the number of UETP passes that you plan to execute. Each pass requires two pages for each line printer and hardcopy terminal.

Check that all terminals are set to the correct baud rate and are assigned appropriate characteristics. (Refer to the user's guide for your terminal.)

Spooled devices and devices allocated to queues fail the initialization phase of UETP and are not tested.

19.3.10 Ethernet Adapters

Make sure that no other processes are sharing the Ethernet adapter device when you run UETP.

Note

UETP will not test your Ethernet adapter if DECnet for OpenVMS or some other application has the device allocated.

Because either DECnet for OpenVMS or the LAT terminal server can try to use the Ethernet adapter (a shareable device), you must shut down DECnet and the LAT terminal server before you run the device test phase, if you want to test the Ethernet adapter.

19.3.11 DR11–W Data Interface (VAX Only)

The DR11–W data interface uses an internal logical loopback mode that tests all features except that of module connectors, cables, and transceivers.

Caution

Only a Compaq support representative can set up the DR11–W data interface for UETP testing.

Because random external patterns are generated during this operation, the user device or other processor might need to be isolated from the DR11–W data interface being tested until the testing is completed.

To test the DR11–W data interface properly, the E105 switchpack must be set as follows:

Switch 1	Switch 2	Switch 3	Switch 4	Switch 5
Off	On	Off	Off	On

When UETP testing is completed, restore the DR11–W data interface to the proper operating configuration.

19.3.12 DRV11–WA Data Interface (VAX Only)

The DRV11–WA data interface is a general-purpose, 16-bit, parallel, direct memory access (DMA) data interface.

Caution

Only a Compaq support representative can set up the DRV11–WA data interface for UETP testing.

To prepare the DRV11–WA driver on a MicroVAX computer for UETP testing, be sure the following conditions exist:

- The jumpers on the DRV11–WA board are set to W2, W3, and W6.
- A loopback cable is connected to the DRV11–WA board.
- The DRV11–WA board occupies slots 8 to 12. If the DRV11–WA is in another location, timeout errors can occur.

When UETP testing is completed, restore the DRV11–WA to the proper operating configuration.

19.3.13 DR750 or DR780 (DR32 Interface) (VAX Only)

The DR32 (DR750 or DR780) device is an interface adapter that connects the internal memory bus of a VAX processor to a user-accessible bus called the DR32 device interconnect (DDI).

Caution

Only a Compaq support representative can set up the DR750 or DR780 for UETP testing.

To prepare the DR750 or the DR780 for UETP testing, use the following procedure:

1. Copy the DR780 microcode file, XF780.ULD, from the diagnostic medium to SYS\$SYSTEM. Use the procedure described in the documentation provided with the DR780 Microcode Kit.
2. Turn off the power to the DR780.
3. Make the following DR780 backplane jumper changes:
 - a. Remove the jumper from W7 and W8.
 - b. Add a jumper from E04M1 to E04R1.
 - c. Add a jumper from E04M2 to E04R2.
4. Disconnect the DDI cable from the DR780. This cable is either a BC06V–nn cable, which can be disconnected, or a BC06R–nn cable, which requires that you remove its paddle card from the backplane of the DR780.
5. Restore power to the DR780.

When UETP testing is completed, restore the DR750 or the DR780 to the proper operating configuration.

Testing the System with UETP

19.3 Setting Up the Devices to Be Tested

19.3.14 Second LPA11–K Device

If you have two LPA11–K devices, be sure that each is given a systemwide logical name in the SYS\$MANAGER:LPA11STRT.COM file. The logical name for the first LPA11–K device should be LPA11\$0, and the logical name for the second LPA11–K device should be LPA11\$1.

19.3.15 Devices That Are Not Tested

UETP does not test the following devices; their status has no effect on UETP execution:

- Devices that require operator interaction (such as card readers)
- Software devices (such as the null device and local memory mailboxes)

UETP does not have specific tests for UDA, HSC, or CI devices; they are tested implicitly by the disk, magnetic tape, and DECnet for OpenVMS tests.

UETP also does not test the console terminal or console drives. If you boot the system, log in, and start UETP, you have shown that these devices can be used.

19.3.16 OpenVMS Cluster Testing

Before you run UETP in an OpenVMS Cluster environment, check the SYSTEST_CLIG account. The SYSTEST_CLIG account parallels SYSTEST except that it is dedicated to running the cluster-integration test. The requirements for the SYSTEST_CLIG account are as follows:

- The account should be present in the user authorization file, exactly as distributed by Compaq on each system in your OpenVMS Cluster.

Note

The SYSTEST_CLIG account could have been disabled during the OpenVMS upgrade procedure. If it was disabled, you must reenab the SYSTEST_CLIG account and give it a null password before you run UETP.

To reenab the SYSTEST_CLIG account, enter the following commands:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN AUTHORIZE
UAF> MODIFY /FLAGS=NODISUSER /NOPASSWORD SYSTEST_CLIG
UAF> EXIT
```

Note

Compaq recommends that you disable the SYSTEST_CLIG account after testing has completed.

To disable the SYSTEST_CLIG account, enter the following commands:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN AUTHORIZE
UAF> MODIFY /FLAGS=DISUSER SYSTEST_CLIG
UAF> EXIT
```

- The privileges and quotas of the SYSTEST_CLIG account must match those of the SYSTEST account.

Testing the System with UETP

19.3 Setting Up the Devices to Be Tested

UETP requires little additional preparation for the cluster-integration test phase beyond the requirements for other UETP test phases. The additional requirements for cluster integration testing are as follows:

1. Your system must be a member of a cluster. If it is not, UETP displays a message and does not attempt to run the test.
2. Your system must use the same deadlock detection interval as the other systems in the cluster. (The deadlock detection interval is set by the system parameter DEADLOCK_WAIT. It is normally not changed from the default value, which is 10 seconds.)
3. The files UETCLIG00.COM and UETCLIG00.EXE, located in SYS\$TEST, are necessary for each system included in the test.
4. DECnet for OpenVMS must be set up between the cluster nodes; UETP uses DECnet for OpenVMS to create a process on those nodes. All checks that the test makes depend on its ability to create the SYSTEST_CLIG processes and to communicate with them using DECnet for OpenVMS software.
5. All operator terminals (OPA0:) should accept broadcast messages. To set the BROADCAST characteristic, enter the following command:

```
$ SET TERM/BROADCAST/PERM OPA0:
```

Nodes on which the operator's terminal (OPA0) is set to the NO BROADCAST terminal characteristic will generate the following error message during the cluster test:

```
*****
* UETCLIG00master *
* Error count = 1 *
*****
-UETP-E-TEXT, 0 operator consoles timed out on the cluster test warning
and 1 operator console rejected it.
-UETP-E-TEXT, Status returned was,
"%SYSTEM-F-DEVOFFLINE, device is not in configuration or not
available"
```

6. There must be a [SYSTEST] or [SYS0.SYSTEST] directory on some disk available to the cluster for each node (both OpenVMS and HSC) in the cluster. The test uses the same directory as the UETP disk test to create a file on each cluster node and to see if some other OpenVMS node in the cluster can share access to that file. There must be one such directory per node; the test continues with the next cluster node once it has finished with a file.
7. By default, the UETP cluster phase selects three nodes from the running cluster for deadlock, disk, and file access testing. However, if you want all cluster nodes tested, enter the following command before invoking UETP:

```
$ DEFINE/GROUP UETP$CTMODE ALL
```

19.3.17 Testing a Small-Disk System

After you install the OpenVMS operating system on a small system disk (for example, an RZ23L), you might not have the 1200 blocks of free disk space required to run UETP successfully. If you do not have 1200 free blocks on your system disk, use VMSTAILOR to remove some files from the system disk before you run UETP. For instructions on using VMSTAILOR, refer to the OpenVMS upgrade and installation manual for your system.

Testing the System with UETP

19.3 Setting Up the Devices to Be Tested

19.3.18 DECnet for OpenVMS Phase

The DECnet for OpenVMS phase of UETP uses more system resources than other tests. You can, however, minimize disruptions to other users by running the test on the least busy node.

By default, the file UETDNET00.COM specifies the node from which the DECnet test will be run. To run the DECnet test on a different node, enter the following command before you invoke UETP:

```
$ DEFINE/GROUP UETP$NODE_ADDRESS node_address
```

This command equates the group logical name UETP\$NODE_ADDRESS to the node address of the node in your area on which you want to run the DECnet phase of UETP.

For example:

```
$ DEFINE/GROUP UETP$NODE_ADDRESS 9.999
```

You can also run the DECnet for OpenVMS test on a different node by entering the following command before you invoke UETP:

```
$ DEFINE/GROUP UETP$NODE_NAME "node" "username password"
```

Note

When you use the logical name UETP\$NODE_ADDRESS, UETP tests only the first active circuit found by NCP (Network Control Program). Otherwise, UETP tests all active testable circuits.

When you run UETP, a router node attempts to establish a connection between your node and the node defined by UETP\$NODE_ADDRESS or UETP\$NODE_NAME. Occasionally, the connection between your node and the router node can be busy or nonexistent. When this happens, the system displays the following error messages:

```
%NCP-F-CONNEC, Unable to connect to listener
-SYSTEM-F-REMRSRC, resources at the remote node were insufficient

%NCP-F-CONNEC, Unable to connect to listener
-SYSTEM-F-NOSUCHNODE, remote node is unknown
```

19.3.19 Vector Processors and the VVIEF (VAX Only)

UETP automatically loads all installed and enabled vector processors during the load phase, and automatically tests all installed and enabled vector processors during the device test phase.

If vector processors are available on the system, check for the VP number by entering the following commands:

```
$ x = F$GETSYI ("VP_NUMBER")
$ SHOW SYMBOL x
```

Multiply the value of x by 3. If the result is greater than the account PRCLM value, then you must increase the SYSTEST account PRCLM quota to match the returned result. For more information, see Chapter 28.

Testing the System with UETP

19.3 Setting Up the Devices to Be Tested

However, UETP cannot load the VAX Vector Instruction Emulation facility (VVIEF) during the load phase, and will not automatically test VVIEF. To test VVIEF, you must perform the following steps before running UETP:

1. Edit the file UETCONT00.DAT to add the following line:

```
Y Y UETVECTOR.EXE "DEVICE_TEST"
```

2. Make sure VVIEF was activated when the system was booted. To determine if the VVIEF was activated, enter the following DCL commands:

```
$ X = F$GETSYI("VECTOR_EMULATOR")
$ SHOW SYMBOL X
```

If the system displays a value of 1, VVIEF is loaded; if the system displays a value of 0, VVIEF is not loaded.

The VVIEF test can be executed as an individual test using the RUN command, as described in Section 19.8.2.

19.4 Starting UETP

When you have logged in and prepared the system and devices, you are ready to begin the test.

To start UETP, enter the following command and press Return:

```
$ @UETP
```

UETP displays the following prompt:

```
Run "ALL" UETP phases or a "SUBSET" [ALL]?
```

Throughout the startup dialog, brackets indicate the default value, which you can choose by pressing Return.

When running UETP for the first time, it is recommended that you choose the default value (ALL) and run all the phases. If you choose ALL, UETP displays three more questions, which are described in Section 19.4.2 through Section 19.4.4. If you want to run all the test phases, skip the next section.

19.4.1 Running a Subset of Phases

You can run a single phase by entering SUBSET or S in response to the following prompt:

```
Run "ALL" UETP phases or a "SUBSET" [ALL]?
```

If you enter S or SUBSET, UETP prompts you for the phase you want to run as follows:

You can choose one or more of the following phases:

```
DEVICE, LOAD, DECNET, CLUSTER
```

Phases(s):

There is no default; enter one or more phase names from the list. Separate two or more phases with spaces or commas.

If your choice includes the LOAD phase, UETP displays three prompts:

```
How many passes of UETP do you wish to run [1]?
How many simulated user loads do you want [n]?
Do you want Long or Short report format [Long]?
```

If you exclude the LOAD phase from your list of choices, UETP responds with only two prompts: the first and the third.

Testing the System with UETP

19.4 Starting UETP

The next three sections discuss how you can respond to these questions. After you have answered the questions, the phase you have selected runs to completion.

19.4.2 Single Run Versus Multiple Passes

If you specified the default ALL or a subset of phases at the last prompt, UETP displays the following message:

```
How many passes of UETP do you wish to run [1]?
```

You can repeat the test run as many times as you want. If you enter 1 in response to the prompt (or press Return for the default), UETP stops after completing a single run. If you specify a number greater than 1, UETP restarts itself until it completes the specified number of passes.

You can run UETP once to check that the system is working, or many times to evaluate the system's response to continuous use. For example, a service technician who is interested only in verifying that a newly installed system works might run UETP once or twice. A manufacturing technician might let the system run for several hours as part of the system integration and test.

When you specify multiple UETP runs, you can request a short console log. (See Section 19.4.4.) Ensure that all line printers and hardcopy terminals have enough paper because each run requires two pages.

19.4.3 Defining User Load for Load Test

After you specify the number of passes, UETP prompts you as follows:

```
How many simulated user loads do you want [n]?
```

Note

UETP displays this prompt only if you choose to run the LOAD phase, either implicitly (by running all phases) or explicitly (by running a subset and specifying the LOAD phase).

The load test simulates a situation in which a number of users (detached processes) are competing for system resources. In response to this prompt, enter the number of users you want to simulate for this test. The number in brackets is the default value that UETP computed for your system. The default value depends on the amount of memory and the paging and swapping space that your system has allocated.

Although the given default value is the best choice, you can increase or decrease the user load by entering your own response to the prompt. However, be aware that an increase can cause the test to fail because of insufficient resources.

If you want to see UETP display the user-load equation as it runs, see Section 19.6.2.

19.4.4 Report Formats

The following prompt allows you to choose between long or short report formats:

```
Do you want Long or Short report format [Long]?
```

19.4.4.1 Long Report Format

If you choose the long report format (the default), UETP sends the following information to the console terminal:

- All error messages
- All output generated at the beginning of all phases and tests
- All output generated at the end of all phases and tests

UETP records all its output in the UETP.LOG file, regardless of your response to this question.

In many cases, it might not be convenient to have UETP write the bulk of its output to the terminal. For example, if you run UETP from a hardcopy terminal, the output printing can slow the progress of the tests. This delay might not be a problem if you have requested only one run; however, you might prefer to use the short format if you intend to run multiple passes of UETP from a hardcopy terminal.

19.4.4.2 Short Report Format

If you request the short format, UETP displays status information at the console, such as error messages and notifications of the beginning and end of each phase. This information enables you to determine whether UETP is proceeding normally. If the short console log indicates a problem, you can look at the file UETP.LOG for further information. UETP.LOG contains all the output generated by the various phases, as well as the status information displayed at the console.

After you choose the report format, UETP initiates its sequence of tests and runs to completion. If UETP does not complete successfully, refer to Section 19.6 for troubleshooting information.

19.5 Stopping a UETP Operation

At the end of a UETP pass, the master command procedure UETP.COM displays the time at which the pass ended. In addition, UETP.COM determines whether UETP needs to be restarted. You can request multiple passes when you start up the test package. (See Section 19.4.2.)

At the end of an entire UETP run, UETP.COM deletes temporary files and does other cleanup activities.

Pressing Ctrl/Y or Ctrl/C lets you terminate a UETP run before it completes normally. Normal completion of a UETP run, however, includes the deletion of miscellaneous files that have been created by UETP for the purpose of testing. Using Ctrl/Y or Ctrl/C can interrupt or prevent these cleanup procedures.

The effect of these control characters depends on what part of UETP you are executing. For an explanation of the organization of UETP and its components, refer to Section 19.8.

19.5.1 Using Ctrl/Y

Press Ctrl/Y to abort a UETP run. Note, however, that cleanup of files and network processes in the [SYSTEST] directory might not be complete.

If you are running an individual test image, pressing Ctrl/Y interrupts the current UETP test and temporarily returns control to the command interpreter. While the test is interrupted, you can enter a subset of DCL commands that are executed within the command interpreter and do not cause the current image to exit.

Testing the System with UETP

19.5 Stopping a UETP Operation

19.5.2 Using DCL Commands

The *OpenVMS User's Manual* contains a table of commands that you can use within the command interpreter. In addition, you can enter any of the following commands:

- The CONTINUE command continues the test from the point of interruption (except during execution of the cluster test).
- The STOP command terminates the test; the test aborts and control returns to the command interpreter.

Note

Using the STOP command can prevent cleanup procedures from executing normally. You should use the EXIT command if you want the image to do cleanup procedures before terminating.

- The EXIT command executes cleanup procedures and terminates the test (except during execution of the cluster test); control returns to the command interpreter.

If you enter any DCL command other than those that execute within the command interpreter, the test does cleanup procedures and terminates, and the DCL command executes.

19.5.3 Using Ctrl/C

Press Ctrl/C to interrupt a UETP run. You cannot continue the same test phase after you press Ctrl/C. UETP automatically goes to the next phase in the master command procedure.

Some UETP phases react to Ctrl/C by cleaning up all activity and terminating immediately. These tests display the following message when they are started:

```
%UETP-I-ABORTC, 'testname' to abort this test, type ^C
```

The phases that do not display the previous message terminate all processes they have started. These processes might not have a chance to complete normal cleanup procedures.

If you are running an individual test image, however, you can use Ctrl/C to terminate the execution of the image and complete cleanup procedures.

Note that Ctrl/C does not complete cleanup procedures for the cluster test.

19.6 Troubleshooting: An Overview

This section explains the role of UETP in interpreting operational errors in an OpenVMS operating system. See Section 19.7 for a discussion of common errors that can appear in a UETP run and describes how to correct them.

19.6.1 Error Logging and Diagnostics

When UETP encounters an error, it reacts like a user program. It either returns an error message and continues, or it reports a fatal error and terminates the image or phase. In either case, UETP assumes the hardware is operating properly and it does not attempt to diagnose the error.

Testing the System with UETP

19.6 Troubleshooting: An Overview

If the cause of an error is not readily apparent, use the following methods to diagnose the error:

- OpenVMS Error Log utility (ERROR LOG)—Run ERROR LOG to obtain a detailed report of hardware and system errors. ERROR LOG reports provide information about the state of the hardware device and I/O request at the time of each error. For information about running ERROR LOG refer to the *OpenVMS System Management Utilities Reference Manual*.
- Diagnostic facilities—Use the diagnostic facilities to test exhaustively a device or medium to isolate the source of the error.

19.6.2 Interpreting UETP Output

You can monitor the progress of UETP tests at the terminal from which they were started. This terminal always displays status information, such as messages that announce the beginning and end of each phase and messages that signal an error.

The tests send other types of output to various log files, depending on how you started the tests. (See Section 19.6.7.) The log files contain output generated by the test procedures. Even if UETP completes successfully, with no errors displayed at the terminal, it is good practice to check these log files for errors. Furthermore, when errors are displayed at the terminal, check the log files for more information about their origin and nature.

Each test returns a final completion status to the test controller image, UETPHAS00, using a termination mailbox. This completion status is an unsigned longword integer denoting a condition value. As a troubleshooting aid, UETPHAS00 displays the test's final completion status using the \$FAO and \$GETMSG system services.

Sometimes, however, the \$FAO service needs additional information that cannot be provided using the termination mailbox. When this happens, UETP displays an error message similar to the following one:

```
UETP-E-ABORT, !AS aborted at !%D
```

When UETP displays these types of error messages, check the log files for more information. You can also run the individual test to attempt to diagnose the problem.

The error messages that appear at the terminal and within the log files have two basic sources:

- UETP tests
- System components that are tested

If you need help interpreting the messages, use the OpenVMS Help Message utility (Help Message) or refer either to the *OpenVMS System Messages and Recovery Procedures Reference Manual*¹ or to the manual that describes the individual system component.

¹ This manual has been archived but is available on the *OpenVMS Documentation CD-ROM*.

Testing the System with UETP

19.6 Troubleshooting: An Overview

19.6.3 Displaying Information on Your Screen

Several parts of UETP, such as some device tests, UETINIT00.EXE, UETCLIG00.EXE, and UETDNET00.COM, let you obtain additional information concerning the progress of the test run or the problems the test encounters. Because this information is usually insignificant, it is not displayed on the screen.

To view the information, enter the following command to define the logical name MODE and run the program:

```
$ DEFINE MODE DUMP
```

19.6.4 Example Screen Display (VAX Only)

The following example shows the output for UETINIT00.EXE on a VAX 6000 computer, and logical name MODE defined as DUMP:

```
$DEFINE MODE DUMP
$RUN UETINIT00

      Welcome to VAX/VMS UETP Version X7.3

%UETP-I-ABORTC, UETINIT00 to abort this test, type ^C

You are running on a VAX 6000-430 CPU with 327680 pages of memory.
The system was booted from _$11$DUA6:[SYS0.].

Run "ALL" UETP phases or a "SUBSET" [ALL]?
How many passes of UETP do you wish to run [1]?

The default number of loads is the minimum result of

1) CPU_SCALE * ((MEM_FREE + MEM_MODIFY) / (WS_SIZE * PER_WS_INUSE))
   7.32 * (( 232390 + 5048) / ( 1024 * 0.20)) = 8486
2) Free process slots = 296
3) Free page file pages / Typical use of page file pages per process
   1099992 / 1000 = 1099

How many simulated user loads do you want [296]?
Do you want Long or Short report format [Long]?

UETP starting at 1-MAR-2001 16:00:43.86 with parameters:
DEVICE LOAD DECNET CLUSTER phases, 1 pass, 296 loads, long report.

$
```

This program does not initiate any phase; it displays the equation used by UETP to determine user load and the specific factors that are employed in the current run.

Respond to the questions by pressing Return. After you respond to the first prompt, the program displays the expressions that determine the default number of simultaneous processes. The following definitions apply:

- CPU_SCALE refers to the relative processing power of the CPU in relation to a VAX 11/780 computer. For example, a VAX 6000-430 computer has a CPU_SCALE of 7.32 because it has 7.32 times the processing power of a VAX 11/780 (1.0) computer.
- MEM_FREE represents memory in pages available to users.
- MEM_MODIFY represents memory pages on the modified page list.
- WS_SIZE represents working set size.

Testing the System with UETP

19.6 Troubleshooting: An Overview

- PER_WS_INUSE represents typical percentage of the working set in active use for each process.

UETINIT00 also displays the specific values represented by the expressions. In this example, UETP selects 296 as the default for simulated user loads, because 296 is the minimum result of the three expressions.

Deassign the logical name MODE before running UETP, unless you prefer to see the user load breakdown every time you run UETP.

19.6.5 Example Screen Display (Alpha Only)

The following example shows the output for UETINIT00.EXE on an Alpha system, with logical name MODE defined as DUMP:

```
$ DEFINE MODE DUMP
$ RUN UETINIT00
```

```
          Welcome to OpenVMS Alpha UETP Version 7.3
%UETP-I-ABORTC, UETINIT00 to abort this test, type ^C
You are running on a AlphaServer 4100 5/533 4MB CPU.
The system was booted from _$4$DKA300:[SYS0.].
Run "ALL" UETP phases or a "SUBSET" [ALL]?
How many passes of UETP do you wish to run [1]?
The default number of loads is the minimum result of
1) (MEM_FREE + MEM_MODIFY) / ( WS_SIZE )
   ( 1807872 +      10496 ) / ( 16512 )                = 110
2) Free process slots                                = 488
3) Free page file pages / Typical use of blocks per process
   650240 / 1000 = 650
How many simulated user loads do you want [110]?
Do you want Long or Short report format [Long]?
UETP starting at 1-MAR-2001 15:53:19.52 with parameters:
DEVICE LOAD DECNET CLUSTER phases, 1 pass, 110 loads, long report.
```

This program does not initiate any phase; it displays the equation used by UETP to determine user load and the specific factors that are employed in the current run.

Respond to the questions by pressing the Return key. After you respond to the first prompt, the program displays the expressions that determine the default number of simultaneous processes. The following definitions apply:

- MEM_FREE represents memory in pagelets available to users.
- MEM_MODIFY represents memory pagelets on the modified page list.
- WS_SIZE represents working set size in pagelets.

UETINIT00 also displays the specific values represented by the expressions. In this example, UETP selects 110 as the default for simulated user loads, because 100 is the minimum result of the three expressions.

Deassign the logical name MODE before running UETP, unless you prefer to see the user load breakdown every time you run UETP.

Testing the System with UETP

19.6 Troubleshooting: An Overview

19.6.6 Defining a Remote Node for UETP Ethernet Testing

Occasionally during the UETUNAS00 test, it is difficult to determine whether the problem reports concern the device under test or the remote device. The easiest way to ensure proper error reporting is to define a *good turnaround*. A good turnaround is a remote node that you know turns around Ethernet packets correctly and is up and waiting in the ready state.

You can make the UETUNAS00 test use a known good turnaround by performing the following actions. In the commands that follow, assume that the *good* device is on node BETA and that node BETA is already defined in the network database.

1. Find the address of the good Ethernet node by using the Network Control Program (NCP). To use NCP, the following conditions must apply:
 - DECnet for OpenVMS must be up and running on the system.
 - The account you are using must have TMPMBX and NETMBX privileges.

Enter the following commands and press Return:

```
$ RUN SYS$SYSTEM:NCP
NCP> TELL BETA SHOW EXECUTOR STATUS
```

If node BETA has not been defined in your network database, NCP displays an error message. In this event, specify another good node and retry the command. Otherwise, see your system or network manager.

NCP displays information similar to the following messages:

```
Node Volatile Status as of 22-JUN-2000 16:13:02
Executor node = 19.007 (BETA)
State = on
Physical address = AA-00-03-00-76-D3
Active links = 6
Delay = 1
```

2. Use the displayed *physical address* (in this case, AA00030076D3) to define the logical name TESTNIADR to point to the good turnaround. Note that you do not specify the hyphens (-).

First, log in to the SYSTEST account. Then enter the following command:

```
$ DEFINE/SYSTEM TESTNIADR AA00030076D3
```

3. Run UETP.
4. When UETP has completed, deassign the logical name TESTNIADR by entering the following command:

```
$ DEASSIGN/SYSTEM TESTNIADR
```

19.6.7 Log Files

UETP stores all information generated by all UETP tests and phases from its current run in one or more UETP.LOG files, and it stores the information from the previous run in one or more OLDUETP.LOG files. If a run of UETP involves multiple passes, there will be one UETP.LOG or one OLDUETP.LOG file for each pass.

At the beginning of a run, UETP deletes all OLDUETP.LOG files, and renames any UETP.LOG files to equivalent versions of OLDUETP.LOG. Then UETP creates a new UETP.LOG file and stores the information from the current pass in it. Subsequent passes of UETP create higher versions of UETP.LOG.

Testing the System with UETP

19.6 Troubleshooting: An Overview

Therefore, at the end of a run of UETP that involves multiple passes, there is one UETP.LOG file for each pass. In producing the files UETP.LOG and OLDUETP.LOG, UETP provides the output from the two most recent runs.

The cluster test creates a NETSERVER.LOG file in SYS\$TEST for each pass on each system included in the run. If the test is unable to report errors (for example, if the connection to another node is lost), the NETSERVER.LOG file on that node contains the result of the test run on that node. UETP does not purge or delete NETSERVER.LOG files; therefore, you must delete them occasionally to recover disk space.

If a UETP run does not complete normally, SYS\$TEST can contain other log files. Ordinarily these log files are concatenated and placed within UETP.LOG. You can use any log files that appear on the system disk for error checking, but you must delete these log files before you run any new tests. You can delete these log files yourself or rerun the entire UETP, which checks for old UETP.LOG files and deletes them.

19.7 Troubleshooting: Possible UETP Errors

This section is intended to help you identify and solve problems you can encounter running UETP. You should refer to this section if you need help understanding a system failure and isolating its cause. This section is not intended as a repair manual and is not expected to diagnose any flaws in your system. It should, however, help you to interpret and act upon the information in the error messages.

If you are unable to correct an error after following the steps in this section, you should contact a Compaq support representative. Any information you can supply about the measures you have taken to isolate the problem will help your a Compaq support representative diagnose the problem.

19.7.1 Summary of Common Failures

The following problems are the most common failures encountered while running UETP:

- Wrong quotas, privileges, or account
- UETINIT01 failure
- UETVECTOR failure (VAX computers only)
- Ethernet device allocated or in use by another application
- Insufficient disk space
- Incorrect cluster setup
- Problems during the load test
- DECnet for OpenVMS error
- Errors logged but not displayed
- No process control block (PCB) or swap slots
- System hangups
- Lack of default access for the file access listener (FAL) object
- Bugchecks and machine checks

Testing the System with UETP

19.7 Troubleshooting: Possible UETP Errors

The sections that follow describe these errors and offer the best course of action for dealing with each one.

19.7.2 Wrong Quotas, Privileges, or Account

If your assigned quotas or privileges do not match standard quotas and privileges for the SYSTEST account, UETP displays the following error message:

```
*****
* UETINIT00 *
* Error count = 1 *
*****
-UETP-W-TEXT, The following:
    OPER privilege,
    BIOLM quota,
    ENQLM quota,
    FILLM quota,
```

are nonstandard for the SYSTEST account and may result in UETP errors.

This message informs you that the OPER privilege and the BIOLM, ENQLM, and FILLM quotas either are not assigned correctly or are not assigned at all.

Note

UETP displays a similar message if you run the cluster integration test phase and the privileges and quotas for the SYSTEST_CLIG account are incorrect. The SYSTEST and SYSTEST_CLIG accounts require the same privileges and quotas. Take the action described in this section for both accounts.

Solution

To correct the problem, use the following procedure:

1. Display all privileges and quotas in effect for the SYSTEST account using the Authorize utility (AUTHORIZE) as follows:

```
$ SET DEFAULT SYS$SYSTEM
$ RUN SYS$SYSTEM:AUTHORIZE
UAF> SHOW SYSTEST
```

Testing the System with UETP

19.7 Troubleshooting: Possible UETP Errors

```

Username: SYSTEST                               Owner: SYSTEST-UETP
Account: SYSTEST                                UIC: [1,7] ([SYSTEST])
CLI: DCL                                         Tables: DCLTABLES
Default: SYS$SYSROOT:[SYSTEST]
LGICMD: LOGIN
Login Flags:
Primary days: Mon Tue Wed Thu Fri Sat Sun
Secondary days:
No access restrictions
Expiration: (none) Pwdminimum: 8 Login Fails: 0
Pwdlifetime: 14 00:00 Pwdchange: 22-JUN-2000 10:12
Last Login: (none) (interactive), (none) (non-interactive)
Maxjobs: 0 Fillm: 100 Bytlim: 65536
Maxacctjobs: 0 Shrfillm: 0 Pbytlim: 0
Maxdetach: 0 BIOLm: 12 JTquota: 1024
Prclm: 12 DIOLm: 55 WSdef: 256
Prio: 4 ASTlm: 100 WSquo: 512
Queprio: 0 TQElm: 20 WSextent: 2048
CPU: (none) Enqlm: 300 Pgflquo: 20480
Authorized Privileges:
  CMKRNL CMEXEC SYSNAM GRPNAM DETACH DIAGNOSE LOG_IO GROUP
  PRMCEB PRMMBX SETPRV TMPMBX NETMBX VOLPRO PHY_IO SYSPRV
Default Privileges:
  CMKRNL CMEXEC SYSNAM GRPNAM DETACH DIAGNOSE LOG_IO GROUP
  PRMCEB PRMMBX SETPRV TMPMBX NETMBX VOLPRO PHY_IO SYSPRV
UAF> SHOW SYSTEST_CLIG
.
.
.
UAF> EXIT

```

2. Make sure the default privileges and quotas assigned to the account match the following list:

Privileges

```

CMKRNL CMEXEC NETMBX DIAGNOSE
DETACH PRMCEB PRMMBX PHY_IO
GRPNAM TMPMBX VOLPRO LOG_IO
SYSNAM SYSPRV SETPRV GROUP

```

Quotas

```

BIOLM: 18 PRCLM: 12
DIOLM: 55 ASTLM: 100
FILLM: 100 BYTLM: 65536
TQELM: 20 CPU: no limit
ENQLM: 300 PGFLQUOTA: 20480
WSDEFAULT: 256 WSQUOTA: 512
WSEXTEXT: 2048

```

3. If any privileges or quotas are incorrect, run `AUTHORIZE` to correct them.

If you are logged in to the wrong account, the following error message asks you to log in to the SYSTEST account:

```

$ @UETP
*****
* UETINIT00 *
* Error count = 1 *
*****
-UETP-E-ABORT, UETINIT00 aborted at 22-JUN-2000 14:24:10.13
-UETP-E-TEXT, You are logged in to the wrong account.
Please log in to the SYSTEST account.
$

```

Testing the System with UETP

19.7 Troubleshooting: Possible UETP Errors

You must run UETP from the SYSTEST account.

19.7.3 UETINIT01 Failure

UETINIT01 failures are related to peripheral devices; this type of error message can indicate any of the following problems:

- Device failure
- Device not supported or not mounted
- Device allocated to another user
- Device write locked
- Lost vacuum on a magnetic tape drive
- Drive off line

In some cases, the corrective action is specified explicitly in the error message. For example, you can receive a message from the operator communication manager (OPCOM) informing you of a problem and recommending a corrective measure:

```
%OPCOM, 22-JUN-2000 14:10:52.96, request 1, from user SYSTEST
Please mount volume UETP in device _MTA0:
%MOUNT-I-OPRQST, Please mount volume UETP in device _MTA0:
```

Other error messages can relate information in which the solution is specified implicitly:

```
%UETP-S-BEGIN, UETDISK00 beginning at 22-JUN-2000 13:34:46.03
*****
* DISK_DRA *
* Error count = 1 *
*****
-UETP-E-TEXT, RMS file error in file DRA0:DRA00.TST
-RMS-E-DNR, device not ready or not mounted
%UETP-S-ENDED, UETDISK00 ended at 22-JUN-2000 13:34:46.80
```

This message tells you that a disk drive is either not ready or not mounted. From this information, you know where to look for the cause of the failure (at the disk drive). If you cannot see the cause of the problem immediately, check the setup instructions in Section 19.3.

In other cases, the cause of a failure might not be obvious from the information in the message. The problem can be related to hardware rather than software. For example, the Ethernet adapter test may produce one of the following messages if UETP does not have exclusive access to the Ethernet adapter:

- Intermodule cable unplugged
- Self-test failure code 0000000

To run the self-test diagnostic on the Ethernet adapter successfully, UETP needs exclusive access to the adapter. As explained in Section 19.3.10, you must shut down DECnet and the LAT terminal server before running the UETP device test phase if you want to test the Ethernet adapter.

Solution

To determine where or when the failure occurs in the execution of UETP, use the following procedure:

Testing the System with UETP

19.7 Troubleshooting: Possible UETP Errors

- Run the device test individually. (See Section 19.4.1.) By doing this, you can determine if the failure can be re-created, and you can isolate the cause of the problem by reproducing it using the least amount of software possible.

For example, if the failure occurs only when you run the entire device phase, and not when you run the affected device test individually, you can conclude the problem is related to device interaction. Conversely, if you can re-create the error by running the single device test, then you have proved that the error is not related to device interaction.

- Run the device test with different media. If your run of the single device test succeeded in reproducing the error, the magnetic tape or disk media could be defective. Running the same test with different media determines whether the original media caused the problem.
- Call a Compaq support representative. If you have tried all the previous steps without solving the problem, you should contact a Compaq support representative.

19.7.4 UETVECTOR Failure (VAX Only)

UETP displays a message similar to the following one to signal a vector processor failure:

```
*****
* UETVECTOR *
* Error count = 1 *
*****
%PPL-S-CREATED SOME, created some of those requested - partial success
-UETP-E-SUBSPNERR, Error spawning subordinate process.
-UETP-E-SCHCTXERR, Error scheduling vector context test subprocess.
-UETP-E-VECCTXERR, Error encountered during vector context testing.
%UETP-I-ENDED, UETVECTOR_0000 ended at 22-JUN-2000 07:37:00.59
```

Solution

See Section 19.3.19 for the correct setup for vector processor testing.

19.7.5 Device Allocated or in Use by Another Application

If DECnet for OpenVMS software or the LAT software is running during the DEVICE phase, the UETUNAS00 test displays the following message:

```
-UETP-W-TEXT, Device is in use by DECnet or another application
```

Other UETP communication device tests display the following message:

```
SYSTEM-W-DEVALLOC, device already allocated to another user
```

Solution

If you want to run the device test on the Ethernet adapter, shut down DECnet and LAT software before beginning the test.

19.7.6 Insufficient Disk Space

When you run continuous passes of UETP, log files accumulate on the disk from which UETP was run. These files reduce the amount of free disk space available for each successive pass. If the amount of disk space available becomes too small for the current load, the following error message appears:

Testing the System with UETP

19.7 Troubleshooting: Possible UETP Errors

```
%UETP-S-BEGIN, UETDISK00 beginning at 22-JUN-2000 08:12:24.34
%UETP-I-ABORTC, DISK_DJA to abort this test, type ^C

*****
* DISK_DJA *
* Error count = 1 *
*****
-UETP-F-TEXT, RMS file error in file DJA0:DJA00.TST
-RMS-F-FUL, device full (insufficient space for allocation)

*****
* DISK_DJA *
* Error count = 2 *
*****
-UETP-F-TEXT, RMS file error in file DJA0:DJA01.TST
-RMS-F-FUL, device full (insufficient space for allocation)
%UETP-E-DESTP, DISK_DJA stopped testing DJA unit 0 at 08:12:36.91
%UETP-S-ENDED, UETDISK00 ended at 22-JUN-2000 08:12:37.98
```

Solution

Make more space available on the disk. You can do this by using one or more of the following techniques:

- Delete unnecessary files to create more space.
- Purge files, if multiple versions exist.
- Mount a volume with sufficient space.
- Check for disk quotas that might be enabled on the disk. If disk quotas are enabled, either disable or increase them. (Refer to the *OpenVMS System Management Utilities Reference Manual* for a description of the Disk Quota utility.)
- Run VMSTAILOR if you have a small-disk system. Refer to the upgrade and installation manual for your operating system for more information.

See Section 19.2.2 and Section 19.3.3 for a further discussion of disk space.

19.7.7 Incorrect Setup of an OpenVMS Cluster System

Most problems that can occur during the cluster-integration test are related to improper setup of the OpenVMS Cluster system or of UETP on the cluster. These problems are most likely to occur at the following stages of the cluster test:

- Near the beginning, when processes on OpenVMS nodes are started
- Toward the end, when cluster file access is checked

The cluster test phase shows that various OpenVMS nodes in your cluster can simultaneously access files on selected nodes in the cluster. First, UETP tries to create a file on a disk drive that is accessible to the other selected nodes in the cluster. The following requirements are for creating a file in the cluster test phase:

- A [SYSTEST] directory must exist on the disk in either the master file directory (MFD) or in the root directory [SYS0.].
- The protection for [SYSTEST] directory must be set to allow the SYSTEST account to create a file in it.

If UETP is unable to find a suitable device on a certain node, the test displays a warning message and proceeds to the next cluster node.

Testing the System with UETP

19.7 Troubleshooting: Possible UETP Errors

Nodes on which the operator's terminal (OPA0) is set to the NO BROADCAST terminal characteristic will generate the following error message during the cluster test:

```
*****
* UETCLIG00master *
* Error count = 1 *
*****
-UETP-E-TEXT, 0 operator consoles timed out on the cluster test warning
  and 1 operator console rejected it.
-UETP-E-TEXT, Status returned was,
  "%SYSTEM-F-DEVOFFLINE, device is not in configuration or not
  available"
```

Disregard this message if OPA0 is set to NO BROADCAST.

Solution

Whenever you suspect a problem, examine the SYS\$TEST:NETSERVER.LOG file that was created when the SYSTEST_CLIG process was created. This file can contain additional error information that could not be transmitted to the node running the test. If it was not possible to create the SYSTEST_CLIG process on some node, the system accounting file for that node might contain a final process status in a process termination record.

The following problems can occur during a cluster test:

- Logging in at other nodes—This problem is due to incorrect setup for the cluster test at the remote OpenVMS node. For example, if you specified a password for the SYSTEST_CLIG account or if you disabled the SYSTEST_CLIG account, the test displays the following message:

```
%SYSTEM-F-INVLOGIN, login information invalid at remote node
```

Refer to Section 19.3.16 and Section 19.6.6 for information about preparing for cluster testing.
- Communicating with other nodes—A message indicates a DECnet problem. Check the NETSERVER.LOG file on the affected node to determine the cause.
- Taking out locks or detecting deadlocks—The most likely cause of this problem is that you are not logged in to the SYSTEST account. Another possibility is that your cluster is not configured properly.
- Creating files on cluster nodes—This problem is due to incorrect setup for the cluster test; refer to Section 19.3.16 for information about preparing for cluster testing.

19.7.8 Problems During the Load Test

A variety of errors can occur during the load test because the command procedures that are started during the tests run several utilities and do many functions. Tracking a problem can be difficult because UETP deletes the log files that are generated during the load test. (See Section 19.8.3.)

Solution

If a problem occurs during the load test and the cause is not obvious, you can modify UETP.COM to preserve the log files as follows:

1. Add the /NODELETE qualifier to the following line:

```
$ TCNTRL UETLOAD00.DAT/PARALLEL_COUNT='LOADS/REPORT_TYPE='REPORT
```

Testing the System with UETP

19.7 Troubleshooting: Possible UETP Errors

2. Delete or comment out the following line:

```
$ DELETE UETLO*.LOG;*
```

Rerun the load test with these changes to try to re-create the problem.

If you re-create the problem, look at the contents of the appropriate log file. You can determine which log file to read by understanding the scheme by which the load test names its processes and log files. (The log file names are derived from the process names.)

The load test creates processes that are named in the following format:

UETLOAD nn_nnnn

For example:

```
%UETP-I-BEGIN, UETLOAD00 beginning at 22-JUN-2000 15:45:08.97
%UETP-I-BEGIN, UETLOAD02_0000 beginning at 22-JUN-2000 15:45:09.42
%UETP-I-BEGIN, UETLOAD03_0001 beginning at 22-JUN-2000 15:45:09.63
%UETP-I-BEGIN, UETLOAD04_0002 beginning at 22-JUN-2000 15:45:10.76
%UETP-I-BEGIN, UETLOAD05_0003 beginning at 22-JUN-2000 15:45:11.28
%UETP-I-BEGIN, UETLOAD06_0004 beginning at 22-JUN-2000 15:45:12.56
%UETP-I-BEGIN, UETLOAD07_0005 beginning at 22-JUN-2000 15:45:13.81
%UETP-I-BEGIN, UETLOAD08_0006 beginning at 22-JUN-2000 15:45:14.95
%UETP-I-BEGIN, UETLOAD09_0007 beginning at 22-JUN-2000 15:45:16.99
%UETP-I-BEGIN, UETLOAD10_0008 beginning at 22-JUN-2000 15:45:19.32
%UETP-I-BEGIN, UETLOAD11_0009 beginning at 22-JUN-2000 15:45:19.95
%UETP-I-BEGIN, UETLOAD02_0010 beginning at 22-JUN-2000 15:45:20.20
%UETP-I-BEGIN, UETLOAD03_0011 beginning at 22-JUN-2000 15:45:21.95
%UETP-I-BEGIN, UETLOAD04_0012 beginning at 22-JUN-2000 15:45:22.99
```

Note that if more than 10 processes are created, the numbering sequence for the UETLOAD nn portion of the process name starts over at UETLOAD02; however, the 4 digits of the $nnnn$ portion continue to increase.

Each load test process creates two log files. The first log file is created by the test controller; the second log file is created by the process itself. The log file to look at for error information about any given load test process is the one that was created by the test controller (the first log file).

The load test log file derives its file name from the process name, appending the last four digits of the process name (from the $nnnn$ portion) to UETLO. The test-controller log file and the process log file for each process use the same file name; however, the process log file has the higher version number of the two. For example, the log files created by the process UETLOAD05_0003 would be named as follows:

UETLO0003.LOG;1 (test-controller log file)

UETLO0003.LOG;2 (process log file)

Make sure that you look at the log file with the lower version number; that file contains the load test commands and error information.

After you have isolated the problem, restore UETP.COM to its original state and delete the log files from the load test (UETLO*.LOG;*); failure to delete these files can result in disk space problems.

19.7.9 DECnet for OpenVMS Error

A DECnet error message can indicate that the network is unavailable.

Solution

- If DECnet for OpenVMS software is included in your system, determine whether the product authorization key (PAK) is registered by entering the following command:

```
$ SHOW LICENSE
```

If the PAK is not registered, invoke the License utility to register it by entering the following command:

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

For information about registering licenses, refer to the following documents:

- The OpenVMS Upgrade and Installation Manual for your operating system
- The *OpenVMS License Management Utility Manual*
- If DECnet for OpenVMS software is not included in your system, ignore the message; it is normal and does not affect the UETP run.

If you encounter other DECnet related errors, you should perform the following actions:

- Run DECnet for OpenVMS software as a single phase (see Section 19.4.1) to determine whether the error can be re-created.
- Use the Help Message or refer to the *OpenVMS System Messages: Companion Guide for Help Message Users*.

19.7.10 Errors Logged but Not Displayed

If no errors are displayed at the console terminal or reported in the UETP.LOG file, you should run ERROR LOG to see if any errors were logged in the ERRLOG.SYS file. Refer to the *OpenVMS System Management Utilities Reference Manual* for information about running the ERROR LOG.

19.7.11 No PCB or Swap Slots

The following error message indicates that no PCB or swap slots are available:

Testing the System with UETP

19.7 Troubleshooting: Possible UETP Errors

```
%UETP-I-BEGIN, UETLOAD00 beginning at 22-JUN-2000 07:47:16.50
%UETP-I-BEGIN, UETLOAD02_0000 beginning at 22-JUN-2000 07:47:16.76
%UETP-I-BEGIN, UETLOAD03_0001 beginning at 22-JUN-2000 07:47:16.92
%UETP-I-BEGIN, UETLOAD04_0002 beginning at 22-JUN-2000 07:47:17.13
%UETP-I-BEGIN, UETLOAD05_0003 beginning at 22-JUN-2000 07:47:17.35
%UETP-I-BEGIN, UETLOAD06_0004 beginning at 22-JUN-2000 07:47:17.61
%UETP-W-TEXT, The process -UETLOAD07_0005- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-TEXT, The process -UETLOAD08_0006- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-TEXT, The process -UETLOAD09_0007- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-TEXT, The process -UETLOAD10_0008- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-TEXT, The process -UETLOAD11_0009- was unable to be created,
the error message is
-SYSTEM-F-NOSLOT, no pcb or swap slot available
%UETP-W-ABORT, UETLOAD00 aborted at 22-JUN-2000 07:47:54.10
-UETP-W-TEXT, Aborted via a user Ctrl/C.
*****
*
*   END OF UETP PASS 1 AT 22-JUN-2000 07:48:03.17
*
*****
```

Solution

To solve this problem, use the following procedure:

1. Individually rerun the phase that caused the error message (the LOAD phase in the previous example) to see if the error can be reproduced.
2. Increase the size of the page file, using either the command procedure `SYS$UPDATE:SWAPFILES.COM` (see Chapter 16) or `SYSGEN` (refer to the *OpenVMS System Management Utilities Reference Manual*).
3. Increase the system parameter `MAXPROCESSCNT`, if necessary.
4. Reboot the system.

19.7.12 No Keyboard Response or System Disk Activity

If the keyboard does not respond or the system disk is inactive, the system might be hung.

Solution

A system hangup can be difficult to trace; you should save the dump file for reference. To learn why the system hung, run the System Dump Analyzer as described in the *OpenVMS VAX System Dump Analyzer Utility Manual* or the *OpenVMS Alpha System Analysis Tools Manual*.

Reasons for a system hangup include the following ones:

- Insufficient pool space—Increase the value of the system parameter `NPAGEVIR` and reboot the system.
- Insufficient page file space—Increase the page file space using the `SYSGEN` as described in the *OpenVMS System Management Utilities Reference Manual*.
- I/O device failure causing driver-permanent loop—Call a Compaq support representative.

19.7.13 Lack of Default Access for the FAL Object

If default FAL access is disabled at the remote node selected by UETP for DECnet testing (the adjacent node on each active circuit, or a node defined by the group logical name UETP\$NODE_ADDRESS), messages similar to the following ones appear:

```
%UETP-W-TEXT, The process -SVA019841 0001- returned a final status of:  
%COPY-E-OPENOUT, error opening !AS aS output
```

These messages are followed by:

```
%COPY-E-OPENOUT, error opening 9999"::SVA019841.D1; as output  
-RMS-E-CRE, ACP file create failed  
-SYSTEM-F-INVLOGIN, login information invalid at remote node  
%COPY-W-NOTCOPIED, SYS$COMMON:[SYSTEST]UETP.COM;2 not copied  
%UETP-E-TEXT, Remote file test data error
```

You can ignore these messages.

19.7.14 Bugchecks and Machine Checks

When the system aborts its run, a bugcheck message appears at the console.

Solution

Call your Compaq support representative. Often a hardware problem causes bugchecks and machine checks; solving bugchecks or machine checks is not easy. However, saving the SYS\$SYSTEM:SYSDUMP.DMP and ERRLOG.SYS files is important so they are available for examination. Knowing whether the failure can be re-created is also important; you can run UETP again to verify the failure.

19.8 UETP Tests and Phases

This section explains, in detail, the organization of UETP and the individual components within the test package. You run UETP by starting a master command procedure containing commands to start each test phase. The procedure begins by prompting you for information needed by the various test phases. (See Section 19.4 for a detailed description of starting UETP.)

The master command procedure, UETP.COM, contains commands that initiate each test phase. UETP.COM also contains commands that do such tasks as defining logical names and manipulating files generated by the tests.

The UETP.COM procedure also issues commands to start the test controlling program UETPHAS00.EXE, which, in turn, controls each test phase. The test controller starts up multiple detached processes. It also reports their completion status and other information the processes report to it.

The sections that follow describe the various UETP test phases.

19.8.1 Initialization Phase

The following actions occur during the initialization phase:

- The image UETINIT00.EXE prompts you for information. (See Section 19.4.) Your information defines variables that affect the execution of UETP tests.
- The image UETINIT01.EXE gathers information about all the controllers in the system and on their associated devices. This image writes the information into a file called UETINIDEV.DAT.

Testing the System with UETP

19.8 UETP Tests and Phases

- Using the information in UETSUPDEV.DAT, UETINIT01.EXE verifies which devices in the configuration are operable by running the appropriate device test. Each device test completes a simple read/write operation to each device. If a device fails this test, the device's entry in UETINIDEV.DAT specifies that the device cannot be tested. As a result, subsequent UETP tests ignore that device.
- For each testable controller, UETINIT01.EXE writes a line into a file called UETCONT00.DAT. The line associates a test file with the controller it tests.

A summary of UETINIDEV.DAT always exists in UETP.LOG, and UETINIT01.EXE sends that summary to the console if you have requested the long report format.

19.8.2 Device Test Phase

The device test phase includes separate tests for each type of device, such as disk, magnetic tape, line printer, and terminal. This section explains the device test phase and presents instructions for testing a single device. If you want to run the entire device test phase individually, refer to Section 19.4.1.

19.8.2.1 How the Device Phase Works

The UETP device test phase starts an executable image, the phase controller UETPHAS00, which creates a detached process for every device controller to be tested. For example, if a system includes three terminal controllers, one line printer controller, and two disk controllers, the image creates six detached processes. In parallel, the detached processes execute images that test the various types of devices.

The initialization phase of UETP creates a file called UETINIDEV.DAT and a file called UETCONT00.DAT. UETINIDEV.DAT contains data on the controllers in the system supported by OpenVMS and their associated devices; UETCONT00.DAT associates a device test image with each testable controller.

UETPHAS00 uses the information in UETCONT00.DAT to find a device controller name to pass to each detached process that it creates. UETPHAS00 passes the controller name by writing it to a mailbox that is SYS\$INPUT to individual tests. Each detached process uses that data to determine which controller to test. The test image then searches UETINIDEV.DAT for the device controller and for all testable units on that controller. The phase controller terminates when all devices on all controllers have completed testing.

Because UETCONT00.DAT is deleted automatically at the end of a UETP run, you cannot run the device phase unless you start UETP.COM; you can run only individual test images. UETINIDEV.DAT exists in SYS\$TEST unless you delete it.

19.8.2.2 Running a Single Device Test

You must be logged in to the SYSTEST account to run the individual tests as described in this section. Also, a copy of UETINIDEV.DAT must exist. If a copy of the file is not present from a previous run (a run of the entire UETP or a run of the device test phase creates UETINIDEV.DAT), you can create it. Note that when you run a single test, no log file is created; the test sends all its output to your terminal.

If you do not want to test all the device types, you can test a specific controller by choosing a test image name from Table 19–1 (for VAX systems) or Table 19–2 (for Alpha systems) and executing it as in the following example:

```
$ RUN UETTTYS00
Controller designation?: TTB
```

UETP prompts you for the controller designation and the device code. Unless you are testing your own terminal, you must explicitly designate a controller name. If you are running the terminal test, you can press Return to test your terminal only.

If you plan to repeat the run several times, you might find it more convenient to define the logical name CTRLNAME as follows:

```
$ DEFINE CTRLNAME TTB
$ RUN UETTTYS00
```

When you define the controller name in this way, the logical name CTRLNAME remains assigned after the test completes. To deassign this logical name, use the DCL command DEASSIGN as follows:

```
$ DEASSIGN CTRLNAME
```

19.8.2.3 Format of UETINIDEV.DAT

The UETINIDEV.DAT file is an ASCII sequential file that you can type or edit if necessary. The contents of this file are shown in the following command sequence:

```
$ TYPE UETINIDEV.DAT

DDB x ddd
UCB y uuuuu nnnnnnnnn.nnn
END OF UETINIDEV.DAT
```

The symbols in this example are defined as follows:

Symbol	Value
x	T, if testable units exist for this controller; N, if this controller is not to be tested
y	T, if this unit is testable; N, if this unit is not testable
ddd	Device controller name, for example DUA
uuuuu	Device unit number, for example 25
nnnnnnnnn.nnn	UETP device test name for the unit, for example, UETDISK00.EXE

UETINIDEV.DAT contains a DDB (device data block) line for each controller connected or visible to your system. After the DDB line is a UCB (unit control block) line for each unit connected to that controller. A device test can test a particular device only if both the DDB line and the UCB line indicate that the device is testable.

19.8.2.4 Running a Test in Loop Mode

If you want to put extra stress on a device, you can run the device test in loop mode, which causes the test to run indefinitely. For example:

```
$ DEFINE MODE LOOP
$ RUN UETDISK00
Controller designation?: DRA
%UETP-I-TEXT, End of pass 1 with 980 iterations at 22-JUN-2000 16:18:51:03
```

Testing the System with UETP

19.8 UETP Tests and Phases

^C

You must use Ctrl/C to terminate the test run. If you use Ctrl/Y, UETP does not complete cleanup procedures.

19.8.2.5 Functions of Individual Device Tests

For each disk in the system, the disk test allocates two files into which it randomly writes blocks of data. The test then checks the data, reports any errors to SYS\$OUTPUT, and deletes the disk files.

When you run the disk test phase in a cluster environment, the test accesses all disks that are mounted by the system being tested, and users of the disk being tested might encounter an insufficient disk space problem. You should warn users on remote nodes (who share disks with users on the local system) that UETP might be testing a disk they are using.

The magnetic tape test exercises all the magnetic tape drives in the system. The test creates a large file on each mounted magnetic tape, into which it writes multiple sequential records of varying sizes. After writing the records, the test rewinds the magnetic tape, validates the written records, and reinitializes the magnetic tape.

The terminal and line printer test generates several pages or screens of output, in which each page or screen contains a header line and a test pattern of ASCII characters. A header line contains the test name, the device name, the date, and the time.

For the laboratory peripheral accelerator (LPA11-K), the test image determines the configuration on the LPA11-K's I/O bus. The image loads all types of microcode to the LPA11-K and reads or writes data for each device on the LPA11-K I/O bus.

The communications device tests fill the transmit message buffer with random data; then, using loopback mode, the tests transmit and receive the message several times. To check that the looped-back data is correct, an AST routine is associated with a \$QIO read to compare the received message against the transmitted message. The procedure is repeated using messages of different lengths.

The interface device tests put the devices they are testing in maintenance mode, write random data, and then verify the data.

The Ethernet adapter test does self-test diagnostics on the device. It also does read and write tasks with test data that uses various adapter modes (such as internal loopback and external loopback).

The vector processor device test performs simple vector-scalar and vector-vector arithmetic operations and compares the results with expected values. The test also uses vector-related system service extensions and forces the system to generate arithmetic and memory management exceptions.

Table 19-1 lists the device test images and the devices to be tested on VAX systems.

Table 19–1 Device Tests (VAX Only)

Test Image Name	Devices Tested
UETDISK00.EXE	Disks
UETTAPE00.EXE	Magnetic tape drives and tape cartridge drives
UETTTYS00.EXE	Terminals and line printers
UETLPAK00.EXE	LPA11–K
UETCOMS00.EXE	DMC11, DMR11
UETDMPF00.EXE	DMF32, DMP11
UETDR1W00.EXE	DR11–W
UETDR7800.EXE	DR780, DR750
UETCDRO00.EXE	RRD40, RRD42, RRD50
UETUNAS00.EXE	Ethernet Adapters
UETVECTOR.EXE	Vector Processor, VVIEF

Table 19–2 lists the device test images and the devices to be tested on Alpha systems.

Table 19–2 Device Tests (Alpha Only)

Test Image Name	Devices Tested
UETDISK00.EXE	Disks
UETTAPE00.EXE	Magnetic tape drives and tape cartridge drives
UETTTYS00.EXE	Terminals and line printers
UETCDRO00.EXE	RRD42
UETUNAS00.EXE	Ethernet adapters

19.8.3 System Load Test Phase

The purpose of the system load test is to simulate a number of terminal users who are demanding system resources simultaneously. The system load tests, directed by the file UETLOAD00.DAT, create a number of detached processes that execute various command procedures. Each process simulates a user logged in at a terminal; the commands within each procedure are the same types of commands that a user enters from a terminal. The load test creates the detached processes in quick succession, and the processes generally execute their command procedures simultaneously. The effect on the system is analogous to an equal number of users concurrently issuing commands from terminals. In this way, the load test creates an environment that is similar to normal system use.

The load test uses the logical name LOADS to determine the number of detached processes to create. When you initiate the UETP command procedure, it prompts for the number of users to be simulated (see Section 19.4.3) and consequently the number of detached processes to be created. Your response, which depends on the amount of memory and the swapping and paging space in your system, defines the group logical name LOADS.

The UETP master command procedure deassigns all group logical names assigned by its tests as part of the termination phase. The group logical name LOADS remains assigned only if the UETP package does not complete normally.

Testing the System with UETP

19.8 UETP Tests and Phases

The command procedures executed by the load test can generate a large amount of output, depending on the number of detached processes created. For each detached process (or user), the test creates a version of an output file called UETLOnnnn.LOG (*nnnn* represents a string of numeric characters). The console displays only status information as the load test progresses.

Whether the load test runs as part of the entire UETP or as an individual phase, UETP combines the UETLOnnnn.LOG files, writes the output to the file UETP.LOG, and deletes the individual output files.

You can run the system load test as a single phase by selecting LOAD from the choices offered in the startup dialog. (See Section 19.4.1.)

19.8.4 DECnet for OpenVMS Test Phase

If DECnet for OpenVMS software is included in your OpenVMS system, a run of the entire UETP automatically tests DECnet hardware and software. Because communications devices are allocated to DECnet and the DECnet devices cannot be tested by the UETP device test, UETP will not test the Ethernet adapter if DECnet for OpenVMS or another application has allocated the device. The DECnet node and circuit counters are zeroed at the beginning of the DECnet test to allow for failure monitoring during the run.

As with other UETP phases, you can run the DECnet for OpenVMS phase individually by following the procedure described in Section 19.4.1.

19.8.4.1 Environment

The DECnet for OpenVMS test will work successfully on OpenVMS systems connected to all DECnet supported node types, including routing and nonrouting nodes and several different types of operating systems (such as RSTS, RSX, TOPS, and RT). To copy files between systems, the remote systems must have some type of default access. The DECnet phase tests the following nodes and circuits:

- The node on which UETP is running.
- All circuits in sequence, unless you have defined the logical name UETP\$NODE_ADDRESS to be the remote node that you want to run the test on. If you have defined a remote node, the DECnet phase tests only one circuit.
- All adjacent or first-hop nodes and all circuits in parallel.

No limit exists on the number of communication lines supported by the tests. A test on one adjacent node should last no more than two minutes at normal communications transfer rates.

Note

UETP assumes your system has default access for the FAL object, even though the network configuration command procedure NETCONFIG.COM does not provide access for the FAL object by default. When you install DECnet software with the defaults presented by NETCONFIG.COM, the UETP DECnet phase can produce error messages. You can ignore these error messages. See Section 19.7.13 for more information.

19.8.4.2 How the DECnet Phase Works

UETP (under the control of UETPHAS00.EXE) reads the file UETDNET00.DAT and completes the following steps during the DECnet for OpenVMS phase:

1. Executes a set of Network Control Program (NCP) LOOP EXECUTOR commands to test the node on which UETP is running.
2. Uses NCP to execute the command SHOW ACTIVE CIRCUITS. The results are placed in UETININET.TMP, from which UETP creates the data file UETININET.DAT. The UETININET.TMP file contains the following information for any circuit in the ON state but not in transition:

- Circuit name
- Node address
- Node name (if one exists)

The UETININET.TMP file is used throughout the DECnet phase to determine which devices to test.

3. Uses the UETININET.TMP file to create an NCP command procedure for each testable circuit. Each command procedure contains a set of NCP commands to zero the circuit and node counters and to test the circuit and adjacent node by copying files back and forth.

Note

If you do not want the counters zeroed, do not test the DECnet for OpenVMS software.

4. Executes the command procedures from Step 3 in parallel to simulate a heavy user load. The simulated user load is the lesser of the following values:
 - The number of testable circuits, multiplied by two
 - The maximum number of user-detached processes that can be created on the system before it runs out of resources (determined by UETINIT00)
5. Executes a program, UETNETS00.EXE, that uses the UETININET.DAT file to check the circuit and node counters for each testable circuit. If a counter indicates possible degradation (by being nonzero), its name and value are reported to the console. All counters are reported in the log file, but only the counters that indicate degradation are reported to the console. An example of UETNETS00 output follows.

```
%UETP-S-BEGIN, UETNETS00 beginning at 22-JUN-2000 13:45:33.18
%UETP-W-TEXT, Circuit DMC-0 to (NODENAME1) OK.
%UETP-I-TEXT, Node (NODENAME2) over DMC-1 response timeouts = 1.
%UETP-I-TEXT, Circuit DMC-1 to (NODENAME2) local buffer errors = 34.
%UETP-I-TEXT, Node (NODENAME3) over DMP-0 response timeouts = 3.
%UETP-S-ENDED, UETNETS00 ended at 22-JUN-2000 13:45:36.34
```

Because degradation is not necessarily an error, the test's success is determined by you, not by the system. The following counters indicate possible degradation:

For Circuits

- Arriving congestion loss
- Corruption loss

Testing the System with UETP

19.8 UETP Tests and Phases

- Transit congestion loss
- Line down
- Initialization failure
- Data errors inbound
- Data errors outbound
- Remote reply timeouts
- Local reply timeouts
- Remote buffer errors
- Local buffer errors
- Selection timeouts
- Remote process errors
- Local process errors
- Locally initiated resets
- Network initiated resets

For Nodes

- Response timeouts
- Received connect resource errors
- Node unreachable packet loss
- Node out of range packet loss
- Oversized packet loss
- Packet format error
- Partial routing update loss
- Verification reject

19.8.5 Cluster-Integration Test Phase

The cluster-integration test phase consists of a single program and a command file that depend heavily on DECnet for OpenVMS software. This phase uses DECnet for OpenVMS software to create SYSTEST_CLIG processes on each OpenVMS node in the cluster and to communicate with each node. SYSTEST_CLIG is an account that is parallel to SYSTEST, but limited so that it can only be used as part of the cluster-integration test. The following restrictions on the SYSTEST_CLIG account are necessary for a correct run of the cluster test phase:

- The account must be enabled and the password must be null. For more information, see Section 19.3.16.
- The UIC must be the same as that of the SYSTEST account.
- The account must have the same privileges and quotas as the SYSTEST account. For more information, see Section 19.7.2.
- The account can allow login only through DECnet for OpenVMS software.
- The account must be locked into running UETCLIG00.COM when it logs in.

Testing the System with UETP

19.8 UETP Tests and Phases

These items are necessary to ensure the security and privacy of your system. If the test cannot create a SYSTEST_CLIG process on an OpenVMS node, it gives the reason for the failure and ignores that node for the lock tests and for sharing access during the file test. Also, the test does not copy log files from any node on which it cannot create the SYSTEST_CLIG process. If a communication problem occurs with a SYSTEST_CLIG process after the process has been created, the test excludes the process from further lock and file sharing tests. At the end of the cluster-integration test, an attempt is made to report any errors seen by that node.

UETCLIG00.EXE has two threads of execution: the primary and the secondary. The first, or primary thread, checks the cluster configuration (OpenVMS nodes, HSC nodes, and the attached disks that are available to the node running the test). For selected OpenVMS nodes, the primary thread attempts to start up a SYSTEST_CLIG process through DECnet software. If the primary thread was able to start a SYSTEST_CLIG process on a node, the node runs the command file UETCLIG00.COM, which starts up UETCLIG00.EXE and runs the secondary execution thread.

The process running the primary thread checks to see that it can communicate with the processes running the secondary threads. It then instructs them to take out locks so that a deadlock situation is created.

The primary thread tries to create a file on some disk on selected OpenVMS and HSC nodes in the cluster. It writes a block, reads it back, and verifies it. Next, it selects one OpenVMS node at random and asks that node to read the block and verify it. The primary thread then extends the file by writing another block and has the secondary thread read and verify the second block. The file is deleted.

The secondary processes exit. They copy the contents of their SYS\$ERROR files to the primary process, so that the UETP log file and console report show all problems in a central place. DECnet for OpenVMS software automatically creates a NETSERVER.LOG in SYS\$TEST as the test is run, so that if necessary, you can read that file later from the node in question.

During the test run, the primary process uses the system service SYS\$BRKTHRU to announce the beginning and ending of the test to each OpenVMS node's console terminal.

You can define the group logical name MODE to the equivalence string DUMP to trace most events as they occur. Note that the logical name definitions apply only to the node on which they were defined. You must define MODE on each node in the cluster on which you want to trace events.

Getting Information About the System

This chapter discusses setting up and maintaining system log files, maintaining error log files, and using system management utilities to monitor the system.

This chapter describes the following tasks:

Task	Section
Using the Error Formatter (ERRFMT)	Section 20.3
Using ERROR LOG to produce reports	Section 20.4
Using DECEvent to report system events	Section 20.5
Setting up, maintaining, and printing the operator log file	Section 20.6
Using security auditing	Section 20.7
Using the Monitor utility to monitor system performance	Section 20.8

This chapter explains the following concepts:

Concept	Section
System log files	Section 20.1
Error logging	Section 20.2
Error Log utility (ERROR LOG)	Section 20.4.1
DECEvent Event Management utility	Section 20.5.1
Operator log file	Section 20.6.1
OPCOM messages	Section 20.6.2
Security auditing	Section 20.7.1
Monitor utility (MONITOR)	Section 20.8.1

20.1 Understanding System Log Files

In maintaining your system, collect and review information about system events. The operating system provides several log files that record information about the use of system resources, error conditions, and other system events. Table 20–1 briefly describes each file and provides references to sections that discuss the files in more detail.

Getting Information About the System

20.1 Understanding System Log Files

Table 20–1 System Log Files

Log File	Description	For More Information
Error log file	The system automatically records device and CPU error messages in this file.	Section 20.2
Operator log file	The operator communication manager (OPCOM) records system events in this file.	Chapter 2 and Section 20.6
Accounting file	The accounting file tracks the use of system resources.	Chapter 21
Security audit log file	The audit server process preallocates disk space to and writes security-relevant system events to this file.	Section 20.7

20.2 Understanding Error Logging

The error logging facility automatically writes error messages to the latest version of the error log file, SYS\$ERRORLOG:ERRLOG.SYS. Error log reports are primarily intended for use by Compaq support representatives to identify hardware problems. System managers often find error log reports useful in identifying recurrent system failures that require outside attention.

Starting with OpenVMS Version 7.2, DECEvent Version 2.9 or later is required for analyzing error log files. DECEvent Version 2.9 provides a separate utility, the Binary Error Log Translation utility, in the DECEvent kit. This utility converts the new Common Event Header (CEH) binary error log file into a binary error log file whose header format and structure can be read by earlier versions of DECEvent and by the older Error Log utility.

For more information about the Binary Error Log Translation utility, refer to its documentation, which is included in the DECEvent kit shipped with the OpenVMS kit.

Parts of the Error Logging Facility

The error logging facility consists of the parts shown in Table 20–2.

Table 20–2 Parts of the Error Logging Facility

Part	Description
Executive routines	Detect errors and events, and write relevant information into error log buffers in memory.
Error Formatter (ERRFMT)	The ERRFMT process , which starts when the system is booted, periodically empties error log buffers, transforms the descriptions of errors into standard formats, and stores formatted information in an error log file on the system disk. (See Section 20.3.2.) The Error Formatter allows you to send mail to the SYSTEM account or another user if the ERRFMT process encounters a fatal error and deletes itself. (See Section 20.3.3.)
Error Log utility (ERROR LOG)	Invokes the Error Log Report Formatter (ERF) , which selectively reports the contents of an error log file. You invoke ERROR LOG by entering the DCL command ANALYZE/ERROR_LOG. (See Section 20.4.2.)

(continued on next page)

Table 20–2 (Cont.) Parts of the Error Logging Facility

Part	Description
DECEvent	Selectively reports the contents of an event log file; you invoke DECEvent by entering the DCL command DIAGNOSE. (See Section 20.5.) DECEvent Version 2.9 and higher includes the Binary Error Log Translation utility.

The executive routines and the Error Formatter (ERRFMT) process operate continuously without user intervention. The routines fill the error log buffers in memory with raw data on every detected error and event. When one of the available buffers becomes full, or when a time allotment expires, ERRFMT automatically writes the buffers to SYS\$ERRORLOG:ERRLOG.SYS.

Sometimes a burst of errors can cause the buffer to fill up before ERRFMT can empty them. You can detect this condition by noting a skip in the error sequence number of the records reported in the error log reports. As soon as ERRFMT frees the buffer space, the executive routines resume preserving error information in the buffers.

The ERRFMT process displays an error message on the system console terminal and stops itself if it encounters excessive errors while writing the error log file. Section 20.3.1 explains how to restart the ERRFMT process.

20.3 Using the Error Formatter

The Error Formatter (ERRFMT) process is started automatically at boot time. The following sections explain how to perform these tasks:

Task	Section
Restart the ERRFMT process, if necessary	Section 20.3.1
Maintain error log files	Section 20.3.2
Send mail if the ERRFMT process is deleted	Section 20.3.3

20.3.1 Restarting the ERRFMT Process

To restart the ERRFMT process, follow these steps:

1. Log in to the system manager's account so that you have the required privileges to perform the operation.
2. Execute the site-independent startup command procedure (STARTUP.COM), specifying ERRFMT as the command parameter, as follows:

```
$ @SYS$SYSTEM:STARTUP ERRFMT
```

Note

If disk quotas are enabled on the system disk, ERRFMT starts only if UIC [1,4] has sufficient quotas.

Getting Information About the System

20.3 Using the Error Formatter

20.3.2 Maintaining Error Log Files

Because the error log file, `SYS$ERRORLOG:ERRLOG.SYS`, is a shared file, `ERRFMT` can write new error log entries while the Error Log utility reads and reports on other entries in the same file.

`ERRLOG.SYS` increases in size and remains on the system disk until you explicitly rename or delete it. Therefore, devise a plan for regular maintenance of the error log file. One method is to rename `ERRLOG.SYS` on a daily basis. If you do this, the system creates a new error log file. You might, for example, rename the current copy of `ERRLOG.SYS` to `ERRLOG.OLD` every morning at 9:00. To free space on the system disk, you can then back up the renamed version of the error log file on a different volume and delete the file from the system disk.

Another method is to keep the error log file on a disk other than the system disk by defining the logical name `SYS$ERRORLOG` to be the device and directory where you want to keep error log files; for example:

```
$ DEFINE/SYSTEM/EXECUTIVE SYS$ERRORLOG DUA2:[ERRORLOG]
```

To define this logical name each time you start up the system, add the logical name definition to your `SYLOGICALS.COM` procedure. See Section 5.2.5 for details.

Be careful not to delete error log files inadvertently. You might also want to adopt a file-naming convention that includes a beginning or ending date for the data in the file name.

20.3.3 Using `ERRFMT` to Send Mail

The Error Formatter (`ERRFMT`) allows users to send mail to the system manager or to another designated user if the `ERRFMT` process encounters a fatal error and deletes itself.

Two system logical names, `ERRFMT$_SEND_MAIL` and `ERRFMT$_SEND_TO`, control this feature:

- `ERRFMT$_SEND_MAIL`

To enable sending mail, must translate to the string `TRUE`, and is case insensitive. Any other value disables the sending of mail.

- `ERRFMT$_SEND_TO`

Must translate to a user name (the current default is `SYSTEM`).

Compaq recommends that you do not use distribution lists and multiple user names.

You can define these logical names in one of two ways:

- Dynamically, using `DCL DEFINE/SYSTEM` commands

After you make the changes, you must stop and restart `ERRFMT` for the changes to take effect.

- Permanently, in `SYS$STARTUP:SYLOGICAL.COM`

The logical names you define take effect the next time the system is rebooted. The following instructions use this method.

Getting Information About the System

20.3 Using the Error Formatter

20.3.3.1 Enabling and Disabling ERRFMT to Send Mail

If `ERRFMT$_SEND_MAIL` is defined to be `TRUE`, you receive a mail message with a subject line saying that `ERRFMT` is about to delete itself. The operator log file and the output displayed at the system console, `OPA0:`, contain more detailed information about the failure encountered and instructions on how to restart `ERRFMT`; however, you are often not at the console to see this information.

If you are using `ERRFMT` in one mode, for example, with sending mail enabled, and you want to disable sending mail, use the system manager's account to edit `SYS$STARTUP:SYLOGICAL.COM`, adding the following command:

```
$ DEFINE/SYSTEM ERRFMT$_SEND_MAIL FALSE
```

To reenable sending mail, use the system manager's account to edit `SYS$STARTUP:SYLOGICAL.COM`, adding the following command:

```
$ DEFINE/SYSTEM ERRFMT$_SEND_MAIL TRUE
```

20.3.3.2 Sending Mail to Another User

Sending mail to the `SYSTEM` account is enabled by default. However, you can define `ERRFMT$_SEND_TO` to send mail to another user if `ERRFMT` is about to delete itself.

To change the user name to receive mail, use the system manager's account to edit `SYS$STARTUP:SYLOGICAL.COM`, adding an appropriate logical name `DEFINE` command. For example:

```
$ DEFINE/SYSTEM ERRFMT$_SEND_TO R_SMITH
```

Compaq recommends that you do not use distribution lists and multiple user names.

20.4 Using ERROR LOG

Use the Error Log utility (`ERROR LOG`) to report selectively on the contents of an error log file. You must have the `SYS$PRV` privilege to run `ERROR LOG`. Starting with OpenVMS Version 7.2, before using the Error Log utility, you need to translate error log files using the Binary Error Log Translation utility, which is part of `DECevent`. For more information, refer to its documentation, which is included in the `DECevent` kit.

20.4.1 Understanding ERROR LOG

`ERROR LOG` supports most OpenVMS supported hardware, such as adapters, disks, tapes, CPUs, and memories, but not all communications devices. Some synchronous communications devices are supported.

The operating system automatically writes messages to the latest version of an error log file, `SYS$ERRORLOG:ERRLOG.SYS`, as the events shown in Table 20–3 occur.

Getting Information About the System

20.4 Using ERROR LOG

Table 20–3 Types of Events Reported in the Error Log File

Event	Description
Errors	Device errors, device timeouts, machine checks, bus errors, memory errors (hard or soft error correcting code [ECC] errors), asynchronous write errors, and undefined interrupts
Volume changes	Volume mounts and dismounts
System events	System startups, messages from the Send Message to Error Logger (\$SNDErr) system service, and time stamps

You can use ERROR LOG to process error log entries for the following forms of optional output:

- Full report of selected entries, which is the default
- Brief report of selected entries
- Summary report of selected entries
- Register dump report of selected device entries
- Binary copy of selected entries
- Binary copy of rejected entries

Section 20.4.2 explains how to produce error log reports. Refer to the *OpenVMS System Management Utilities Reference Manual* for examples of error log reports.

The error reports that ERROR LOG produces are useful in two ways:

- They aid preventive maintenance by identifying areas within the system that show potential for failure.
- They aid the diagnosis of a failure by documenting the errors and events that led up to it.

The detailed contents of the reports are most meaningful to Compaq support representatives. However, you can use the reports as an important indicator of the system's reliability. For example, using the DCL command SHOW ERROR, you might see that a particular device is producing a relatively high number of errors. You can then use ERROR LOG to obtain a more detailed report and decide whether to consult your support representative.

If a system component fails, a Compaq support representative can study the error reports of the system activity leading up to and including the failure. If a device fails, you can generate error reports immediately after the failure; for example:

- One report might describe in detail all errors associated with the device that occurred within the last 24 hours.
- Another report might summarize all types of errors for all devices that occurred within the same time period.
- The summary report can put the device errors into a systemwide context.

Your support representative can then run the appropriate diagnostic program for a thorough analysis of the failed device. Using the combined error logging and diagnostic information, your support representative can proceed to correct the device.

Error reports allow you to anticipate potential failures. Effective use of the Error Log utility in conjunction with diagnostic programs can significantly reduce the amount of system downtime.

20.4.2 Producing Error Log Reports

You enter the DCL command in the following format:

```
ANALYZE/ERROR_LOG [/qualifier(s)][filespec[,...]]
```

where:

<i>qualifier</i>	Specifies the function the ANALYZE/ERROR_LOG command is to perform.
<i>filespec</i>	Specifies one or more files that contain information to be interpreted for the error log report.

Refer to the *OpenVMS System Management Utilities Reference Manual* for details about the command and its parameters and for examples of error log reports.

ERROR LOG issues error messages for inconsistent error log entries. Use the Help Message facility to look up explanations and suggested user actions for these messages.

20.4.3 Producing a Full Error Log Report

The following steps show how to produce an error log report for all entries in the error log file and how to print the report:

1. Either log in to the SYSTEM account or ensure that you have the SYSPRV privilege. (You must have privilege to access the error log file.) For example:

```
$ SET PROCESS/PRIVILEGE=SYSPRV
```
2. Set your default disk and directory to SYS\$ERRORLOG:

```
$ SET DEFAULT SYS$ERRORLOG
```
3. Examine the error log directory to see which error log file you want to analyze:

```
$ DIRECTORY
```
4. To obtain a full report of the current error log file, enter the following command:

```
$ ANALYZE/ERROR_LOG/OUTPUT=ERRORS.LIS
```
5. Print a copy of the report, using the file name you specified with the /OUTPUT qualifier:

```
$ PRINT ERRORS.LIS
```

Example

```
$ SET PROCESS/PRIVILEGE=SYSPRV
$ SET DEFAULT SYS$ERRORLOG
$ DIRECTORY ❶
Directory SYS$SYSROOT:[SYSERR]
ERRLOG.OLD;2 ERRLOG.OLD;1 ERRLOG.SYS;1
Total of 3 files.
$ ANALYZE/ERROR_LOG/OUTPUT=ERRORS.LIS ERRLOG.OLD ❷
$ PRINT ERRORS.LIS ❸
```

Getting Information About the System

20.4 Using ERROR LOG

The following list explains the commands in the example.

- ❶ The `DIRECTORY` command lists all the files in the `SYS$ERRORLOG` directory. The directory contains three files: two old error log files and the current error log file, `ERRLOG.SYS`.
- ❷ The `ANALYZE/ERROR_LOG` command writes a full report to a file called `ERRORS.LIS`, using the most recent `ERRLOG.OLD` file as input.
- ❸ The `PRINT` command prints `ERRORS.LIS`.

20.4.4 Using Other Error Log Report Options

This section briefly explains how to specify report formats and produce a report of selected entries.

Table 20–4 contains error log report options. For more details about options and examples of error log reports using options, refer to the *OpenVMS System Management Utilities Reference Manual*.

Table 20–4 Error Log Report Options

In Order To...	You Can...
Specify report formats	Change report formats by using qualifiers, including the following ones: <ul style="list-style-type: none">• <code>/BINARY</code>—to convert binary error log records to ASCII text or to copy error log records to a specified output file.• <code>/BRIEF</code>—to create a brief report.• <code>/REGISTER_DUMP</code>—to generate, in a hexadecimal longword format, a report that consists of device register information (used in conjunction with the <code>/INCLUDE</code> qualifier).• <code>/REJECTED</code>—to specify the name of a file that will contain binary records for rejected entries.
Specify a display device for reports	Use the <code>/OUTPUT</code> qualifier to send reports to a terminal for display or to a disk or magnetic tape file. By default, the system sends the report to the <code>SYS\$OUTPUT</code> device. Because error log reports are 72 columns wide, you can display them on the terminal screen.
Produce a report of selected entries	Use qualifiers to produce error log reports for specific types of events and for a specified time interval. For example, you can process error log entries by selecting a time interval using the <code>/SINCE</code> , <code>/BEFORE</code> , or <code>/ENTRY</code> qualifiers. You can specify error log entries for specific events by using the qualifiers <code>/INCLUDE</code> and <code>/EXCLUDE</code> . These qualifiers form a filter to determine which error log entries are selected or rejected. In addition, you can generate error log reports for one or more OpenVMS Cluster members by using the <code>/NODE</code> qualifier.
Exclude unknown error log entries	By default, when <code>ANALYZE/ERROR_LOG</code> encounters an unknown device, CPU, or error log entry, the utility produces the entry in hexadecimal longword format. Exclude these entries from the report by specifying <code>/EXCLUDE=UNKNOWN_ENTRIES</code> in the command line.

20.5 Using DECEvent

The DECEvent Event Management utility (DECEvent) provides an interface between a system user and the operating system's event log files.

Note

On Alpha DS, ES, and GS systems (other than the AlphaServer GS60 and GS140 systems) running OpenVMS, use the Web-Based Enterprise Services (WEBES), which includes the Compaq Analyze, Compaq Crash Analysis Tool (CCAT), and the Revision and Configuration Management (RCM) tools. You can find WEBES and its documentation on the Compaq System Tools CD-ROM, which is included in the OpenVMS Version 7.3 CD-ROM package. You can find the most up-to-date service tool information at the following location:

<http://www.support.compaq.com/svctools>

You cannot use WEBES on the AlphaServer GS60 or the AlphaServer GS140. DECEvent and the WEBES tools can be used together in a cluster.

20.5.1 Understanding DECEvent

DECEvent allows system users to produce ASCII reports derived from system event entries. The format of the ASCII reports depends on the command entered on the command language interpreter (CLI) with a maximum character limit of 255 characters.

DECEvent uses the error log file, SYS\$ERRORLOG:ERRLOG.SYS, as the default input file, unless you specify another input file.

Event reports are useful for determining preventive maintenance by helping to identify areas within the system showing potential failure. Event reports also aid in the diagnosis of a failure by documenting events that led to the failure.

The contents of the event reports are most meaningful to your Compaq support representative. However, you can use the event reports as an indicator of system reliability. For example, while using the DCL command SHOW ERROR, you might see that a particular device is producing a higher than normal number of events. You can use DECEvent to obtain various detailed reports and determine if you need to contact your Compaq support representative.

If a system component fails, your Compaq support representative can use the event reports to create a history of events leading up to and including the failure.

Used in conjunction with diagnostic programs, event reports significantly reduce the amount of system down time.

DECEvent Report Types

DECEvent produces five types of reports:

Report Type	Description
Full (default)	Provides a translation of all available information for each entry in the event log.
Brief	Provides a translation of key information for each entry in the event log.
Terse	Provides binary event information and displays register values and other ASCII messages in a condensed format.
Summary	Provides a statistical summary of the event entries in the event log.

Getting Information About the System

20.5 Using DECEvent

Report Type	Description
Fast Error (FSTERR)	Provides a quick, one-line per-entry report of your event log for a variety of disk devices.

These report types are mutually exclusive; in other words, you can select only one report type in a command.

Section 20.5.5 contains examples of types of reports. The *OpenVMS System Management Utilities Reference Manual* contains additional examples of the types of reports produced by DECEvent.

The following sections explain how to use DECEvent:

Task	Section
Invoking and exiting DECEvent	Section 20.5.2
Using DECEvent qualifiers	Section 20.5.3
Using additional DECEvent commands	Section 20.5.4
Producing DECEvent reports	Section 20.5.5

In addition, restrictions are listed in Section 20.5.6.

20.5.2 Invoking and Exiting DECEvent

To invoke DECEvent, enter the following DCL command:

```
DIAGNOSE [/primary qualifier][/secondary qualifier[,...][file-spec][,...]
```

Note that you do not need to enter the /TRANSLATE qualifier on the command line because it is the default primary qualifier.

A brief discussion of valid qualifiers, their uses, and their order is provided in the following sections. For a more detailed discussion, see the *DECEvent User's Guide*. This guide is available online at the following web site:

<http://www.support.compaq.com/svctools>

Select “support tools”, next “DECEvent”, and then “documentation”. The *DECEvent User's Guide* is available in several formats.

To exit DECEvent, press Ctrl/C and the Return key (to display the system prompt).

You must have SYSPRV privilege to run DECEvent; however, only read access is required to access the ERRLOG.SYS file. You must have the DIAGNOSE privilege for the /CONTINUOUS primary qualifier to work, enabling the continuous display of events on a terminal screen.

20.5.3 Using DECEvent Qualifiers

The DECEvent qualifiers shown and described in Table 20–5 and Table 20–6 allow you to change the format of the reports that DECEvent produces.

Getting Information About the System

20.5 Using DECEvent

Table 20–5 Primary Qualifiers

Qualifier	Description
/ANALYZE	Provides analysis of the event log or real-time analysis of the event logging utilities.
/BINARY	Controls whether the binary error log records are converted to ASCII text or copied to the specified output file. Do not use this qualifier with any report type qualifier (/FULL, /BRIEF, /TERSE, /SUMMARY, and /FSTERR) or with the /OUTPUT qualifier.
/CONTINUOUS	Specifies that events are formatted in real time, as they are logged by the operating system event logger.
/DUMP	Specifies the output to be a brief report followed by a dump of information from the input event log file.
/INTERACTIVE	Allows users to exit from the command line interface and enter the DECEvent interactive command shell.
/TRANSLATE (default)	Provides translation of event log files into reports.

Table 20–6 (Optional) Secondary Qualifiers

Qualifier	Description
/BEFORE	Specifies that only those entries dated earlier than the stated date and time are to be selected for the event report.
/BRIEF	Generates a brief report
/ENTRY	Generates a report that includes the specified entry range or starts at the specified entry number.
/EXCLUDE	Excludes events generated by the specified device class, device name, or error log entry type from the report.
/FSTERR	Generates a quick, one-line-per-entry report of an event log entry for disks and tapes.
/FULL (default)	Generates a full report, which provides all available information about an event log entry.
/INCLUDE	Includes events generated by the specified device class, device name, or error log entry type in the report.
/LOG	Controls whether informational messages that specify the number of entries selected and rejected for each input file are sent to SYS\$OUTPUT.
/NODE	Generates a report consisting of event entries for specific nodes in an OpenVMS Cluster system.
/OUTPUT	Specifies the output file for the report.
/REJECTED	Allows you to specify the name of a file that will contain binary records for rejected entries.
/SINCE	Specifies that only those entries dated later than the stated date and time are to be selected for the report.
/SUMMARY	Generates an event report that consists of a statistical summary.
/TERSE	Generates an event report consisting of binary event information, register values, and ASCII messages in a condensed format.

Do not use the /BINARY qualifier with any report type qualifier (/FULL, /BRIEF, /TERSE, /SUMMARY, and /FSTERR) or with the /OUTPUT qualifier.

Getting Information About the System

20.5 Using DECEvent

Privileges Required

- You must have SYSPRV privilege to run DECEvent; however, only read access is required to access the ERRLOG.SYS file.
- You must have the DIAGNOSE privilege for the /CONTINUOUS qualifier to work, enabling the continuous display of events on a terminal screen.

20.5.4 Using Additional DECEvent Commands

In addition to the qualifiers listed in Table 20–5 and Table 20–6, DECEvent contains a set of DIRECTORY commands and a set of SHOW commands:

- DIRECTORY commands
These commands allow you to display a list of rulesets that DECEvent needs to translate events into a readable format. (A **ruleset** is a software routine or function that is analogous to an executable file.)
The following DIRECTORY commands are currently implemented in DECEvent:
 - DIRECTORY EVENT
This command lists all rulesets associated with event translation.
 - DIRECTORY CANONICAL
This command lists all rulesets associated with event reports.
- SHOW commands
These commands allow a user to view specific settings and selections. The following SHOW commands are currently implemented in DECEvent:
 - SHOW SELECT
By appending a specific selection keyword name to the SHOW SELECT command, you view only that selection keyword.
 - SHOW SETTINGS
By appending a specific setting's name to the SHOW SETTINGS command, you view only that setting's name and value.

20.5.5 Producing DECEvent Reports

This section contains examples of DECEvent commands and reports.

20.5.5.1 Producing a Full Report

To produce a full report, use the /FULL qualifier. The full report format provides a translation of all available information for each entry in the event log. The full report is the default report type if a report type is not specified in the command line.

Both of the following commands will produce a full report format:

```
$ DIAGNOSE/FULL
```

```
$ DIAGNOSE
```

(/FULL is the default.)

Example 20–1 shows the format of a full report.

Getting Information About the System 20.5 Using DECEvent

Example 20-1 Full Report Format

```
***** ENTRY 1 *****
Logging OS 1. OpenVMS
System Architecture 2. Alpha
OS version V7.3
Event sequence number 1583.
Timestamp of occurrence 18-APR-2000 09:21:18
System uptime in seconds 58004.
Error mask x00000000
Flags x0001 Dynamic Device Recognition present
Host name COGENT

Alpha HW model DEC 3000 Model 400
System type register x00000004 DEC 3000
Unique CPU ID x00000002
mpnum x000000FF
mperr x000000FF

Event validity -1. Unknown validity code
Event severity -1. Unknown severity code
Entry type 100.
Major Event class 3. IO Subsystem

IO Minor Class 1. MSCP
IO Minor Sub Class 5. Logged Message

---- Device Profile ----
Vendor
Product Name RAID 0 - Host Based
Unit Name COGENT$DPA
Unit Number 10.
Device Class x0001 Disk

---- IO SW Profile ----
VMS DC$_CLASS 1.
VMS DT$_TYPE 175.

---- MSCP Logged Msg ----

Logged Message Type Code 22. RAID Message
RAID Event Type 8. Remove Member
Distinguished Member 0.
Member Index 1.
RAID Urgency 4. Global Disk Error
RAID Status x00180009 Bit 00 - Reduced
Bit 03 - Striped
Bit 19 - FE Dis FE
Bit 20 - BC Buff Copy Off

RAIDset Name KGB
*****
```

20.5.5.2 Producing a Brief Report

To produce a brief report, use the /BRIEF qualifier. The brief report format provides translation of key information for each entry in the event log. For example:

```
$ DIAGNOSE/BRIEF
```

Getting Information About the System

20.5 Using DECEvent

Example 20–2 shows the format of a brief report.

Example 20–2 Brief Report Format

```
***** ENTRY      1 *****
Logging OS                1. OpenVMS
System Architecture       2. Alpha
OS version                V7.3
Event sequence number    1583.
Timestamp of occurrence  18-APR-2000 09:21:18
System uptime in seconds 58004.
Error mask               x00000000
Host name                COGENT
Alpha HW model           DEC 3000 Model 400
System type register     x00000004 DEC 3000
Unique CPU ID           x00000002
mpnum                   x000000FF
mperr                   x000000FF
Event validity           -1. Unknown validity code
Event severity           -1. Unknown severity code
Major Event class       3. IO Subsystem
IO Minor Class          1. MSCP
IO Minor Sub Class      5. Logged Message
---- Device Profile ----
Vendor
Product Name            RAID 0 - Host Based
Unit Name               COGENT$DPA
Unit Number             10.
Device Class            x0001 Disk
Logged Message Type Code 22. RAID Message
RAID Event Type         8. Remove Member
Distinguished Member    0.
Member Index            1.
RAID Urgency            4. Global Disk Error
RAID Status             x00180009 Bit 00 - Reduced
                        Bit 03 - Striped
                        Bit 19 - FE Dis FE
                        Bit 20 - BC Buff Copy Off
RAIDset Name            KGB
*****
```

20.5.5.3 Producing a Terse Report

To produce a terse report, use the /TERSE qualifier. The terse report format provides binary event information and displays register values and other ASCII messages in a condensed format. For example:

```
$ DIAGNOSE/TERSE
```

Example 20–3 shows the format of a terse report.

Getting Information About the System 20.5 Using DECEvent

Example 20–3 Terse Report Format

```
***** ENTRY 1 *****
Logging OS 1.
System Architecture 2.
OS version V7.3
Event sequence number 1583.
Timestamp of occurrence 2000041809211800
System uptime in seconds 58004.
Error mask x00000000
Flags x0001
Host name COGENT

Alpha HW model DEC 3000 Model 400
System type register x00000004
Unique CPU ID x00000002
mpnum x000000FF
mperr x000000FF

Event validity -1.
Event severity -1.
Entry type 100.
Major Event class 3.

IO Minor Class 1.
IO Minor Sub Class 5.

---- Device Profile ----
Vendor
Product Name RAID 0 - Host Based
Unit Name COGENT$DPA
Unit Number 10.
Device Class x0001

---- IO SW Profile ----
VMS DC$ CLASS 1.
VMS DT$ TYPE 175.

---- MSCP Logged Msg ----
Logged Message Type Code 22.
RAID Event Type 8.
Distinguished Member 0.
Member Index 1.
RAID Urgency 4.
RAID Status x00180009
RAIDset Name KGB

*****
```

20.5.5.4 Producing a Summary Report

To produce a summary report, use the /SUMMARY qualifier. The summary report format provides a statistical summary of the event entries in the event log. For example:

```
$ DIAGNOSE/SUMMARY
```

Example 20–4 shows the format of a summary report.

Getting Information About the System

20.5 Using DECEvent

Example 20-4 Summary Report Format

```
SUMMARY OF ALL ENTRIES LOGGED ON NODE COGENT

IO Subsystem
  MSCP                      9.
  Host Based RAID          3.

DATE OF EARLIEST ENTRY      18-APR-2000 09:21:18
DATE OF LATEST ENTRY        12-MAY-2000 10:44:54
```

20.5.5.5 Producing a Fast Error (FSTERR) Report

To produce a Fast Error report, use the /FSTERR qualifier. For example:

```
$ DIAGNOSE/FSTERR
```

The Fast Error report provides a quick, one-line-per-entry report of your event log for a variety of disk devices. This makes event analysis and system troubleshooting much easier by eliminating extraneous event information. For example:

```
$ DIAGNOSE/FSTERR [infile]
```

A Fast Error report is shown in Example 20-5.

Example 20-5 Fast Error (FSTERR) Report Format

Drive Name	yymmdd	hhmmss	Entry	MSCP Evtnt	LED	LBN	Physical Cyl Hd Sec	RA	RP	Drive/ HSC Volume Serial
LUKE\$DUA070	921119	160754	3	00EB		255			70 71	V00717
LUKE\$DUA070	921119	160754	4	00EB		255			70 71	V00717
HSC015\$DUA028	910323	113204	5	00EB					70 51	V15039
HSC015\$DUA028	910323	113204	6	00EB					71 51	V15039
BATES\$DUA197	921118	002116	7	00EB					72 32	V17524
CHEWIE\$DUA101	911205	114908	8	00EB					73 81	V 17
PMASON\$DUA006	921207	165007	15	00EB		255			90 42	D23387
PMASON\$DUA006	921207	165007	16	00EB		255			90 42	D23387
C3P0\$DUA242	870218	060031	17	01AB					90 40	D48575
CHER\$DU2132*901008	231053		18	00EB					92 81	D 2345

The Fast Error report includes the information needed by a Compaq support representative to troubleshoot a problem with a tape or disk device.

20.5.6 DECEvent Restrictions

When you use the DECEvent utility, note some of the restrictions listed in this section.

Page File Quota

Sometimes, if the page file quota is exceeded, DECEvent will terminate and return you to the system prompt. If this happens, invoke the last command.

Logical File Names

DECEvent does not translate as input any logical defined as a search list of file names. For example:


```
$ DEFINE EVENT LOG DISK1:[EVENTS]EVENT_LOG1.SYS,DISK1:EVENT_LOG.SYS
$ DIAGNOSE/ANALYZE EVENT_LOG

DECEvent T1.0 FT2
  _DIAGNOSE-FAT: Analyze - No files found ' event_log '
  _DIAGNOSE-FAT: An error occurred while executing a command ruleset
  _DIAGNOSE-INF: No Error Messages to send in thread 1
```

Log File Purging

DECEvent does not automatically purge log files. Set the version limit on the files and directory to your preference. For example:

```
$ SET FILE/VERSION=3 DIA_ACTIVITY.LOG
```

System-Initiated Call Logging

When a system running DECEvent is shut down and rebooted, DECEVENT\$STARTUP.COM does not define FMGPROFILE logicals. This can interfere with proper logging of system initiated call logging (SICL) due to missing customer profile information in the SICL message text.

Unrecognized Messages

The DIAGNOSE command does not recognize error log messages logged using the \$SNDERR system service.

20.6 Setting Up, Maintaining, and Printing the Operator Log File

The following sections describe the contents of the operator log file and OPCOM messages. They also explain how to perform the following tasks, which require OPER privilege:

Task	Section
Setting up the operator log file	Section 20.6.3
Maintaining the operator log file	Section 20.6.4
Printing the operator log file	Section 20.6.5

20.6.1 Understanding the Operator Log File

The operator log file (SYS\$MANAGER:OPERATOR.LOG) records system events and user requests that the operator communication manager (OPCOM) sends to the operator terminal. This recording occurs even if all operator terminals have been disabled. By default, OPCOM starts when you boot your system. (For more information about OPCOM, see Section 2.4.)

You can use the operator log file to anticipate and prevent hardware and software failures and to monitor user requests for disk and magnetic tape operations. By regularly examining the operator log file, you can often detect potential problems and take corrective action.

The size of and access to the OPERATOR.LOG file (or to the file pointed to by the logical OPC\$LOGFILE_NAME) is limited by the size and access of the disk device on which it resides. If disk device does not have enough room to write to the log file or if access to the device in any other way is restricted, records might be missing from the log file.

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

20.6.2 Understanding OPCOM Messages

The following sections describe the types of messages that appear in the operator log file.

Type of Message	Section
Initialization messages	Section 20.6.2.1
Device status messages	Section 20.6.2.2
Terminal enable and disable messages	Section 20.6.2.3
User request and operator reply messages	Section 20.6.2.4
Volume mount and dismount messages	Section 20.6.2.5
System parameter messages	Section 20.6.2.6
Security alarm messages	Section 20.6.2.7

Section 20.6.2.8 contains an example of typical kinds of messages found in an operator log file.

20.6.2.1 Initialization Messages

When you enter the REPLY/LOG command, the system closes the current operator log file and creates and opens a new version of the file. The system records all subsequent OPCOM messages in the new log file.

When you create a new log file, the first message recorded in it is an initialization message. This message shows the terminal name of the operator who initialized the log file, and the log file specification. This message appears in the following format:

```
%%%%%%%%%%  %OPCOM, <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%%  
Logfile has been initialized by operator <terminal-name>  
Logfile is <logfile-specification>
```

Example

```
%%%%%%%%%%  OPCOM, 19-APR-2000 12:29:24.52 %%%%%%%%%%%  
Logfile has been initialized by operator MARS$VTA2:  
Logfile is HOMER::SYS$SYSMOND:[SYSMGT]OPERATOR.LOG;43
```

20.6.2.2 Device Status Messages

Some I/O drivers send messages to OPCOM concerning changes in the status of the devices they control. For example, when a line printer goes off line, an OPCOM message appears in the operator log file at periodic intervals until you explicitly return the device to online status.

The device status message appears in the operator log file in the following format:

```
%%%%%%%%%%  OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%%  
Device <device-name> is offline
```

This message can appear for card readers, line printers, and magnetic tapes.

20.6.2.3 Terminal Enable and Disable Messages

The following sections explain commands you can give to enable and disable terminals as operator terminals (or consoles) and explanations of the corresponding messages that appear in the operator log file.

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

REPLY/ENABLE Messages

To designate a terminal as an operator terminal, enter the REPLY/ENABLE command from the desired terminal. OPCOM confirms the request by displaying messages in the following format at the operator terminal and in the operator log file:

```
%%%%%%%%%% %OPCOM dd-mmm-yyyy hh:mm:ss.cc %%%%%%%%%%%  
Operator <terminal-name> has been enabled, username <user-name>  
%%%%%%%%%% %OPCOM dd-mmm-yyyy hh:mm:ss.cc %%%%%%%%%%%  
Operator status for operator <terminal-name>  
<status-report>
```

These messages tell you which terminal has been established as an operator terminal and lists the requests the terminal can receive and respond to.

You can also designate a terminal as an operator terminal for a particular function by entering the REPLY/ENABLE=*class* command.

If you enter the command REPLY/ENABLE=TAPES, for example, OPCOM displays messages similar to the following ones:

```
%%%%%%%%%% %OPCOM 19-APR-2000 10:25:35.74 %%%%%%%%%%%  
Operator _ROUND$OPA1: has been enabled, username SYSTEM  
%%%%%%%%%% %OPCOM 19-APR-2000 10:25:38.82 %%%%%%%%%%%  
Operator status for operator _ROUND$OPA1:  
TAPES
```

OPCOM confirms that the terminal is established as an operator terminal and indicates that the terminal can only receive and respond to requests concerning magnetic-tape-oriented events, such as the mounting and dismounting of tapes.

REPLY/DISABLE Messages

A terminal that you designate as an operator terminal automatically returns to nonoperator status when the operator logs out. To return the terminal to normal (nonoperator) status without logging out, enter the REPLY/DISABLE command from the terminal.

OPCOM confirms that the terminal is no longer an operator terminal by displaying a message both at the operator terminal and in the operator log file. The message, which tells you which terminal has been restored to nonoperator status and when the transition occurred, has the following format:

```
%%%%%%%%%% %OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%%  
Operator <terminal-name> has been disabled, username <user-name>
```

If you designate a terminal as an operator terminal and only partial operator status is disabled, OPCOM displays a status message. This message lists which requests the terminal can still receive and respond to. This message is displayed at the operator terminal and in the operator log file in the following format:

```
%%%%%%%%%% %OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%%  
Operator status for operator <terminal-name>  
<status-report>
```

For example, suppose you designate a terminal as an operator terminal that receives messages concerning magnetic tapes and disks, as well as messages intended for the special site-specific operator class known as OPER10. Later, you relinquish the terminal's ability to receive messages concerning tapes. When you enter the REPLY/DISABLE=TAPES command, OPCOM returns a message like the following one:

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

```
%%%%%%%%%% %Opcom 19-APR-2000 09:23:45.32 %%%%%%%%%%
Operator status for operator TTA3
DISKS, OPER10
```

This message tells you that terminal TTA3 still receives and can respond to messages about disks and messages directed to OPER10.

20.6.2.4 User Request and Operator Reply Messages

To communicate with the operator, the user enters the REQUEST command, specifying either the /REPLY or /TO qualifier. The following table contains explanations of these qualifiers:

Command	Explanation
REQUEST/REPLY	The request is recorded in the operator log file in the following format: <pre>%%%%%%%%%% %OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%% Request <request-id>, from user <user-name> on <node-name> <_terminal-name:>, <"message-text"></pre> This message tells you which user sent the message, the time the message was sent, the request identification number assigned to the message, the originating terminal, and the message itself.
REQUEST/TO	The request is recorded in the operator log file in the format shown in the REQUEST/REPLY example, but without a request identification number: <pre>%%%%%%%%%% %OPCOM, <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%% Request from user <user-name> on <node-name> <_terminal-name:>, <"message-text"></pre>

Messages also differ depending on how you reply to a user:

Command	Explanation
REPLY/TO	The response is recorded in the operator log file in the following format: <pre>response message <hh:mm:ss.cc>, request <request-id> completed by operator <terminal-name></pre> This message indicates how the operator responded to the user's request, as well as when the response was entered and which operator responded.
REPLY/ABORT	The response is recorded in the operator log file in the following format: <pre><hh:mm:ss.cc>, request <request-id> was aborted by operator <terminal-name></pre>
REPLY/PENDING	The response is not recorded in the operator log file because the request has not yet been completed (that is, the request has not been fulfilled or aborted).

When a user enters a REQUEST/REPLY command and you have disabled all terminals as operators' terminals, OPCOM records all subsequent users' requests in the log file, but returns a message to the user indicating that no operator coverage is available.

All other OPCOM responses to REPLY commands, except responses involving the REPLY/ENABLE, REPLY/DISABLE, and REPLY/LOG commands, are not logged in the operator log file.

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

20.6.2.5 Volume Mount and Dismount Messages

Perhaps the widest range of operator messages occurs with volume mounts and dismounts; for example:

```
%%%%%%%%%% OPCOM, 19-APR-2000 22:41:07.54 %%%%%%%%%%%  
message from user SYSTEM  
Volume "KLATU" " " dismounted, on physical device MTA0:  
15-APR-2000 22:42:14.81, request 2 completed by operator OPA0
```

20.6.2.6 System Parameter Messages

Users with the appropriate privileges can change the following sets of values for system parameters:

Values	Description
Current	Values stored in the default parameter file on disk and used to boot the system
Active	Values stored in memory and used while the system is running

When the system boots, it reads the current values into memory, creating active values. An active value remains equal to the current value until you change either value.

Users can make the following changes to active and current system parameters:

- Active system parameters—Users with CMKRNL privilege can use the System Management utility (SYSMAN) or the System Generation utility (SYSGEN) to change system parameters in the running (active) system. Users can change only those active values that are categorized as *dynamic* system parameters.
- Current system parameters—Users with SYSPRV privilege can use SYSMAN or SYSGEN to change system parameters in the current system.

Note

Compaq recommends that you use AUTOGEN or SYSMAN, not SYSGEN, to change system parameters, as explained in Section 15.2.

OPCOM logs all changes made to active and current system parameters with messages in the following format:

```
%%%%%%%%%% %OPCOM <dd-mmm-yyyy hh:mm:ss.cc> %%%%%%%%%%%  
Message from user <user-name>  
%SYSGEN-I-WRITExxx, <system-mode> system parameters modified by process ID  
<process-id> into file <file-spec>
```

Example

```
%%%%%%%%%% %OPCOM 3-JUN-2000 08:11:59.55 %%%%%%%%%%%  
Message from user D_PLUTO on ANASAT  
%SYSGEN-I-WRITECUR, CURRENT system parameters modified by process ID 000020B  
into file SYS$UPDATE:[SYSTEM]UPDATESYS.PAR;2
```

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

This message indicates that current system parameters have been changed.

Note

If you have changed the format of system messages with the DCL command SET MESSAGE, these messages might not appear in the log file.

20.6.2.7 Security Alarm Messages

Alarm messages are sent to the security operator terminal when selected events occur. See Section 20.7.6 for instructions about how to enable a terminal to receive security alarm messages.

Example

The following example shows a security alarm OPCOM message after a change to JTQUOTA:

```
%%%%%%%% OPCOM 6-JAN-2000 10:41:21.10 %%%%%%%%%
Message from user AUDIT$SERVER on BISCO
Security alarm (SECURITY) and security audit (SECURITY) on BISCO, system id:
20353
Auditable event:      System UAF record modification
Event time:          6-JAN-2000 10:41:20.69
PID:                 00600123
Process name:        SYSTEM
Username:            SYSTEM
Process owner:       [SYSTEM]
Terminal name:       RTA1:
Image name:          BISCO$DUA0:[SYS0.SYSCOMMON.][SYSEXE]AUTHORIZE.EXE
Object class name:   FILE
Object name:         SYS$SYSTEM:SYSUAF.DAT;4
User record:         NEWPORT
JTQUOTA:             New:           2048
                    Original:      1024
```

20.6.2.8 Contents of an Operator Log File

Example 20–6 illustrates some typical messages found in an operator log file.

Example 20–6 Sample Operator Log File (SYS\$MANAGER:OPERATOR.LOG)

```
%%%%%%%% OPCOM, 19-APR-2000 22:26:07.90 %%%%%%%%%
Device DMA0: is offline. ❶
Mount verification in progress.
%%%%%%%% OPCOM, 19-APR-2000 22:26:20.22 %%%%%%%%%
Mount verification completed for device DMA0:
%%%%%%%% OPCOM, 19-APR-2000 22:33:54.07 %%%%%%%%%
Operator ' ZEUS$VT333:' has been disabled, user JONES ❷
%%%%%%%% OPCOM, 19-APR-2000 22:34:15.47 %%%%%%%%%
Operator ' ZEUS$VT333:' has been enabled, user SMITH
%%%%%%%% OPCOM, 19-APR-2000 22:34:15.57 %%%%%%%%%
operator status for ' ZEUS$VT333:'
PRINTER, TAPES, DISKS, DEVICES
%%%%%%%% OPCOM, 19-APR-2000 22:38:53.21 %%%%%%%%%
request 1, from user PUBLIC ❸
Please mount volume KLATU in device MTA0:
The tape is in cabinet A
```

(continued on next page)

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

Example 20–6 (Cont.) Sample Operator Log File (SYS\$MANAGER:OPERATOR.LOG)

```
%%%%%%%%%% OPCOM, 19-APR-2000 22:39:54.37 %%%%%%%%%%
request 1 was satisfied.
%%%%%%%%%% OPCOM, 19-APR-2000 22:40:23.54 %%%%%%%%%%
message from user SYSTEM ④
Volume "KLATU      " mounted, on physical device MTA0:
%%%%%%%%%% OPCOM, 19-APR-2000 22:40:38.02 %%%%%%%%%%
request 2, from user PUBLIC
MOUNT new relative volume 2 ( ) on MTA0:
%%%%%%%%%% OPCOM, 19-APR-2000 22:41:07.54 %%%%%%%%%%
message from user SYSTEM
Volume "KLATU      " dismounted, on physical device MTA0:
15-APR-2000 22:42:14.81, request 2 completed by operator OPA0
%%%%%%%%%% OPCOM, 19-APR-2000 22:46:47.96 %%%%%%%%%%
request 4, from user PUBLIC
_TTB5:, This is a sample user request with reply expected.
%%%%%%%%%% OPCOM, 19-APR-2000 22:47:38.50 %%%%%%%%%%
request 4 was canceled
%%%%%%%%%% OPCOM, 19-APR-2000 22:48:21.15 %%%%%%%%%%
message from user PUBLIC
_TTB5:, This is a sample user request without a reply expected.
%%%%%%%%%% OPCOM, 19-APR-2000 22:49:37.64 %%%%%%%%%%
Device DMA0: has been write locked.
Mount verification in progress.
%%%%%%%%%% OPCOM, 19-APR-2000 23:33:54.07 %%%%%%%%%%
message from user NETACP
DECnet shutting down
```

The following messages appear in the example:

- ① Device status message
- ② Terminal enable and disable message
- ③ User request and operator reply messages
- ④ Volume mount and dismount messages

20.6.3 Setting Up the Operator Log File

The operator log file normally resides on the system disk in the [SYSMGR] directory. You can, however, maintain the log file in a different location by defining the logical name OPC\$LOGFILE_NAME.

The size of and access to the OPERATOR.LOG file (or to the file pointed to by the logical OPC\$LOGFILE_NAME) is limited by the size and access of the disk device on which it resides. If the disk device does not have enough room to write to the log file or if access to the device is restricted in any other way, records might be missing from the log file.

Because this file is in ASCII format, you can print it. Print copies regularly and retain these copies for reference. Section 20.6.5 describes how to print copies of the operator log file.

The system creates a new version of OPERATOR.LOG each time the system is rebooted (except on workstations in an OpenVMS Cluster environment, where the log file is not opened by default). Note that one operator log file exists for each node; it is not a shared file.

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

20.6.3.1 Creating a New Version of the Operator Log File

You can use the DCL command `REPLY/LOG` to create a new version of the file at any time. The highest version is always the one in use and is inaccessible to other users. By default, messages of all operator classes are in the log file.

The following list contains guidelines for using the `REPLY/LOG` command:

- You can use the `REPLY/LOG/ENABLE=(keyword)` and `REPLY/LOG/DISABLE=(keyword)` commands to specify which operator classes to include in the log file.
- When you use the `/LOG` qualifier with the `REPLY/ENABLE` and `REPLY/DISABLE` commands, the classes you select are enabled or disabled for the log file rather than for the terminal.

If a log file is already open, the list of classes is preserved and enabled on the newly created log file. If a log file is not open, the value of the logical `OPC$ENABLE_LOGFILE_CLASSES` is used. If that logical does not exist, all classes are enabled on the new log file.

For more information, refer to the `REPLY/LOG`, `REPLY/ENABLE`, and `REPLY/DISABLE` commands in the *OpenVMS DCL Dictionary*.

Example

The following command opens a log file to include messages about mounting and dismounting disks and tapes:

```
$ REPLY/LOG/ENABLE=(DISKS,TAPES)
```

20.6.3.2 Specifying Logical Names

You can specify the default state of the operator log files by defining logical names in the command procedure `SYS$MANAGER:SYLOGICALS.COM`. The following table lists these logical names and their functions. For more information about `SYLOGICALS.COM`, see Section 5.2.5.

Caution

Setting the `OPC$ALLOW_INBOUND` and `OPC$ALLOW_OUTBOUND` logical names to `FALSE` severs all OPCOM traffic in the specified direction. All OPCOM messages, as well as any returned status messages that might be expected, will not be delivered.

Logical Name	Function
<code>OPC\$ALLOW_INBOUND</code>	Allows OPCOM traffic that is inbound to the node to be turned on or off. By default, this logical name is set to <code>TRUE</code> . If this logical name is set to <code>FALSE</code> , the node will not receive any OPCOM messages from other nodes in the cluster.
<code>OPC\$ALLOW_OUTBOUND</code>	Allows OPCOM traffic that is outbound from the node to be turned on or off. By default, this logical name is set to <code>TRUE</code> . If this logical name is set to <code>FALSE</code> , the node will not send any OPCOM messages to other nodes in the cluster.
<code>OPC\$LOGFILE_ENABLE</code>	Specifies whether an operator log file is opened. If defined to be true, an operator log file is opened. If defined to be false, no operator log file is opened. By default, a log file is opened on all systems except workstations in an OpenVMS Cluster environment.

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

Logical Name	Function
OPC\$LOGFILE_CLASSES	Specifies the operator classes that are enabled for the log file. By default, a log file is opened for all classes. The logical name can be a search list of the allowed classes, a comma-separated list, or a combination of the two. Note that you can define OPC\$LOGFILE_CLASSES even if you do not define OPC\$LOGFILE_ENABLE. In that case, the classes are used for any log files that are opened, but the default is used to determine whether to open the log file.
OPC\$LOGFILE_NAME	Specifies the name of the log file. By default, the log file is named SYS\$MANAGER:OPERATOR.LOG. If you specify a disk other than the system disk, include commands to mount that disk in the command procedure SYLOGICALS.COM.
OPC\$OPA0_ENABLE	Overrides values of symbols for workstations in a cluster. If you define the logical as TRUE, it sets the OPA0 device to BROADCAST (overrides the NOBROADCAST default setting). For systems that are not workstations in a cluster, if you define the logical as FALSE, it sets the OPA0 device to NOBROADCAST.

Note

The only logical that is used for more than the initial startup of OPCOM is OPC\$LOGFILE_NAME. All other OPCOM logicals are ignored. For example, a REPLY/LOG command opens a new operator log file even if the logical OPC\$LOGFILE_ENABLE is defined to be false. To reset OPCOM states and classes after startup, use REPLY/ENABLE or REPLY/DISABLE commands.

20.6.4 Maintaining the Operator Log File

Devise a plan for regular maintenance of operator log files. One way is to start a new log file and rename the second-highest version daily. (See the example in the next section.) You might want to purge outdated versions of the operator log file on a regular basis. However, do not delete versions that you have not backed up. For more information, see Section 5.2.7.9.

If OPCOM is inadvertently deleted, follow these steps to start it manually:

1. Log in to the SYSTEM account so that you have the required privileges to perform the operation.
2. Enter the following command to execute the startup command procedure (STARTUP.COM), specifying OPCOM as the command parameter:

```
$ @SYS$SYSTEM:STARTUP OPCOM
```

20.6.5 Printing the Operator Log File

Perform the following operation to produce a printed copy of the most recent version of the operator log file. (You must have OPER privilege.)

1. Use the following command to enable the terminal as an operator terminal:

```
$ REPLY/ENABLE
```

2. Close the current log file and open a new one by entering the following command:

```
$ REPLY/LOG
```

Getting Information About the System

20.6 Setting Up, Maintaining, and Printing the Operator Log File

3. Set the default to SYS\$MANAGER and enter the following command to list all versions of the file:

```
$ SET DEFAULT SYS$MANAGER
$ DIRECTORY OPERATOR.LOG
```

4. Rename the second-highest version to OPERATOR.OLD:

```
$ RENAME OPERATOR.LOG;-1 OPERATOR.OLD
```

The version number, `-1`, specifies that you want to rename the second-highest version of this file. (The highest version number is the current operator log file.)

5. Print the operator log file by entering the following command:

```
$ PRINT OPERATOR.OLD
```

Example

```
$ REPLY/ENABLE ❶
$ REPLY/LOG ❷

%%%%%%%%%% OPCOM, 19-APR-2000 12:28:20.11 %%%%%%%%%%%
Logfile was closed by operator MARS$VTA2: ❸
Logfile was HOMER::SYS$MANAGER:[SYSMGT]OPERATOR.LOG;27
%%%%%%%%%% OPCOM, 19-APR-2000 12:29:24.52 %%%%%%%%%%%
Logfile has been initialized by operator MARS$VTA2:
Logfile is HOMER::SYS$MANAGER:[SYSMGT]OPERATOR.LOG;28

$ SET DEFAULT SYS$MANAGER ❹
$ DIRECTORY OPERATOR.LOG ❺

Directory SYS$MANAGER:[SYSMGT]
OPERATOR.LOG;28          OPERATOR.LOG;27

Total of 2 files.

$ RENAME OPERATOR.LOG;-1 OPERATOR.OLD ❻
$ PRINT OPERATOR.OLD ❼
```

The following list provides explanations of the numbered commands and responses in the example:

- ❶ The `REPLY/ENABLE` command enables the terminal as an operator terminal.
- ❷ The `REPLY/LOG` command closes the current log file and opens a new one.
- ❸ The response from OPCOM verifies that it has opened a new log file.
- ❹ The `SET DEFAULT` command sets the operator default disk to the system disk.
- ❺ The `DIRECTORY` command displays the files in the directory `[SYSMGT]` on the system disk.
- ❻ The `RENAME` command renames the second-highest version of the operator log file to `OPERATOR.OLD`.
- ❼ The `PRINT` command prints the old operator log file, `OPERATOR.OLD`.

20.7 Using Security Auditing

This section discusses how security auditing works; it also explains how to enable security auditing and how to create a new version of the security audit log file. For more information about the security audit log file, refer to the *OpenVMS Guide to System Security*.

20.7.1 Understanding Security Auditing

Security auditing is the act of recording security-relevant events as they occur on the system. Security-relevant events are divided into a number of categories called **event classes**.

By default, the system enables security auditing when you install or upgrade your system for the events shown in Table 20–7.

Table 20–7 Event Classes Audited by Default

Class	Description
ACL	Access to any object holding a security Auditing ACE.
AUDIT	All uses of the SET AUDIT command. You cannot disable this category.
AUTHORIZATION	All changes to the authorization database: <ul style="list-style-type: none">• System user authorization file (SYSUAF)• Network proxy authorization files: NETPROXY and NET\$PROXY• Rights database (RIGHTSLIST)
BREAKIN	All break-in attempts: batch, detached, dialup, local, network, remote.
LOGFAILURE	All login failures: batch, dialup, local, remote, network, subprocess, detached.

If the security requirements at your site justify additional auditing, you can enable security auditing for other event classes by using the DCL command SET AUDIT, as explained in Section 20.7.4.

20.7.1.1 Security Audit Log File

The audit server process, which is created at system startup, records the events that are shown in Table 20–7 in the security audit log file, SYS\$MANAGER:SECURITY.AUDIT\$JOURNAL.

The usefulness of the security audit log file depends upon the procedures you adopt to review the file on a regular basis. For example, you might implement the following procedure as part of your site audit review policy:

1. Create a new version of the security audit log file each morning.
2. Review the previous version of the log file for suspicious system activity. Depending on the number of security events you are auditing on your system, it might be impractical to review every audit record written to the audit log file. In that case, you might want to select a specific set of records from the log file (for example, all Authorization and Breakin records, or all events created outside normal business hours).
3. If, during your review, you find any security events that appear suspicious, perform a more detailed inspection of the security audit log file, as described in the *OpenVMS Guide to System Security*.

Getting Information About the System

20.7 Using Security Auditing

20.7.1.2 Audit Log Files in Mixed-Version Clusters

The Audit Analysis utility (ANALYZE/AUDIT) running on earlier-version systems is unable to process the current version of audit log files. You must use the current version of ANALYZE/AUDIT to process the current version of the audit log files. The recommended procedure is to maintain separate audit log files on mixed-version clusters.

If redirecting the audit log files, issue the following command on both the earlier-version node and on the node running the current version:

```
AUDIT/JOURNAL/DESTINATION=filespec
```

The destination filespec is stored in the audit server database file. By default, the files are stored in SYS\$COMMON:[SYSMGR] and are called SECURITY_AUDIT.AUDIT\$JOURNAL and SECURITY.AUDIT\$JOURNAL, respectively.

The operating system allows workstations and other users with limited management resources to duplicate their audit log files on another node. The secondary log, a security archive file, is then available to a security administrator on a remote node who has the skills to analyze the file.

Each node in a cluster must have its own archive file. An archive file cannot be shared by multiple nodes in a cluster.

Refer to the *OpenVMS Guide to System Security* for more information.

20.7.2 Displaying Security Auditing Information

To see which event classes your site currently audits, you can enter the DCL command SHOW AUDIT.

The follow example contains security information:

```
$ SHOW AUDIT

System security alarms currently enabled for:
ACL
Breakin:      dialup,local,remote,network,detached
Privilege use:
SECURITY
Privilege failure:
SECURITY

System security audits currently enabled for:
ACL
Authorization
Breakin:      dialup,local,remote,network,detached
Login:       dialup,local,remote,network,detached
Logfailure:   batch,dialup,local,remote,network,subprocess,detached
Logout:      dialup,local,remote,network,detached
Privilege use:
SECURITY
Privilege failure:
ACNT      ALLSPOOL  ALTPRI    AUDIT     BUGCHK    BYPASS    CMEXEC    CMKRNL
DETACH    DIAGNOSE  EXQUOTA  GROUP     GRPNAM    GRPPRV    LOG_IO    MOUNT
NETMBX    OPER      PFNMAP    PHY IO    PRMCEB    PRMGBL    PRMMBX    PSWAPM
READALL   SECURITY  SETPRV    SHARE     SHMEM     SYSGBL    SYSLCK    SYSNAM
SYSPRV    TMPMBX    VOLPRO    WORLD

DEVICE access:
Failure:     read,write,physical,logical,control
FILE access:
Failure:     read,write,execute,delete,control
VOLUME access:
Failure:     read,write,create,delete,control
```

20.7.3 Delaying Startup of Auditing

Ordinarily, the system turns on auditing in VMS\$LPBEGIN just before SYSTARTUP_VMS.COM executes. You can change this behavior, however, by redefining the logical name SYS\$AUDIT_SERVER_INHIBIT.

To change the point at which the operating system begins to deliver security-event messages, add the following line to the SYS\$MANAGER:SYLOGICALS.COM command procedure:

```
$ DEFINE/SYSTEM/EXECUTIVE SYS$AUDIT_SERVER_INHIBIT YES
```

You can initiate auditing during another phase of system startup, perhaps at the end of SYSTARTUP_VMS.COM, by editing the command file to add the following line:

```
$ SET AUDIT/SERVER=INITIATE
```

For information about editing SYSTARTUP_VMS.COM, see Section 5.2.7.

20.7.4 Enabling Security Auditing for Additional Classes

To enable security auditing for classes in addition to those shown in Table 20–7, use the following format:

```
SET AUDIT/ENABLE=keyword[,...] {/ALARM | /AUDIT}
```

The *OpenVMS Guide to System Security* contains descriptions of event classes that you can enable.

When you enable auditing for additional event classes, you must specify two qualifiers:

1. /ENABLE
2. Either /ALARM or /AUDIT (Although you *must* specify one qualifier, you *can* specify both.)

The following table contains explanations of the /ENABLE, /ALARM, and /AUDIT qualifiers.

Qualifier	Explanation
/ENABLE	Defines which event classes you want audited. See Chapter 21 for more information.
/ALARM /AUDIT	<p>Defines the destination of the event message.</p> <ul style="list-style-type: none"> • /ALARM directs the message to all enabled security operator terminals. • /AUDIT directs the message to the security audit log file. <p>Use the /ALARM and /AUDIT qualifiers to report critical events. Less critical events can be written only to the security audit log file for later examination.</p> <p>The default event classes listed in Table 20–7 are sent as both alarms and audits.</p>

The system begins auditing new events on all nodes as soon as you enable them.

Getting Information About the System

20.7 Using Security Auditing

Examples

1. The command in the following example enables auditing for volume mounts and dismounts and sends messages to the security audit log file.

```
$ SET AUDIT/ENABLE=MOUNT/AUDIT
```

2. The command in the following example enables auditing of unsuccessful file accesses and sends messages to all enabled security operator terminals as well as to the security audit log file.

```
$ SET AUDIT/ALARM/AUDIT/ENABLE=ACCESS=FAILURE/CLASS=FILE
```

20.7.5 Disabling Security Auditing

The system continues auditing until you explicitly disable the classes with the /DISABLE qualifier using the following syntax:

```
SET AUDIT/DISABLE=keyword[...] {/ALARM | /AUDIT}
```

20.7.6 Enabling a Terminal to Receive Alarm Messages

The system sends alarm messages to terminals enabled for security class messages. Security alarm messages are not written to the operator log file. They appear only on terminals enabled for security class messages.

In most cases, security alarm messages appear on the system console by default. Since messages scroll quickly off the screen, it is good practice to enable a separate terminal for security class messages and disable message delivery to the system console.

Either choose a terminal in a secure location that provides hardcopy output, or have dedicated staff to monitor the security operator terminal. You can enable any number of terminals as security operators.

To set up a terminal to receive security class alarms, enter the following DCL command from the designated terminal:

```
$ REPLY/ENABLE=SECURITY
```

Example

The following example shows a security alarm message:

```
%%%%%%%%%% OPCOM 25-MAY-2000 16:07:09.20 %%%%%%%%%%%
Message from user AUDIT$SERVER on GILMORE
Security alarm (SECURITY) on GILMORE, system id: 20300
Auditable event:      Process suspended ($SUSPND)
Event time:           25-MAY-2000 16:07:08.77
PID:                  30C00119
Process name:         Hobbit
Username:             HUBERT
Process owner:        [LEGAL,HUBERT]
Terminal name:        RTA1:
Image name:           $99$DUA0:[SYS0.SYSCOMMON.][SYSEXE]SET.EXE
Status:               %SYSTEM-S-NORMAL, normal successful completion
Target PID:           30C00126
Target process name:  SMISERVER
Target username:      SYSTEM
Target process owner: [SYSTEM]
```

20.7.7 Generating Security Reports

The most common type of report to generate is a brief, daily listing of events. You can create a command procedure that runs in a batch job every evening before midnight to generate a report of the day's security event messages and send it to the system manager via Mail.

Note

Because the MOUNT command translates /NOLABEL to /FOREIGN in the audit record, use ANALYZE/AUDIT/SELECT=MOUNT_FLAGS=FOREIGN instead of ANALYZE/AUDIT/SELECT=MOUNT_FLAGS=NOLABEL.

The following example shows the ANALYZE/AUDIT command line you would use to generate this type of report:

```
$ ANALYZE/AUDIT/SINCE=TODAY/OUTPUT=31JAN2000.AUDIT -  
$ SYS$MANAGER:SECURITY.AUDIT$JOURNAL  
$ MAIL/SUBJECT="Security Events" 31JAN2000.AUDIT SYSTEM
```

20.7.8 Creating a New Version of the Security Audit Log File

Because the security audit log file continues to grow until you take action, you must devise a plan for maintaining it.

Use the SET AUDIT command to create a new version of the clusterwide security audit log file. To prevent the loss of audit messages, the previous version of the audit log file is not closed until all audit messages stored in memory are written to the file.

20.7.8.1 Creating a New Clusterwide Version of the Log File

To open a new, clusterwide version of the security audit log file, use the following command:

```
$ SET AUDIT/SERVER=NEW_LOG
```

The audit server process opens a new version of the audit log file on each cluster node.

After you open the new log, rename the old version, using a naming convention for your files that incorporates in the file name a beginning or ending date for the data. Then copy the file to another disk, delete the log from the system disk to save space, and run the Audit Analysis utility on the old log.

By archiving this file, you maintain a clusterwide history of auditing messages. If you ever discover a security threat on the system, you can analyze the archived log files for a trail of suspicious user activity during a specified period of time.

20.7.8.2 Creating a New Node-Specific Version of the Log File

In some cases, OpenVMS Cluster nodes might not share the same system security audit log file. To create a new, node-specific version of the security audit log file, use the following commands:

```
$ SET AUDIT/DESTINATION=filespec  
$ SET AUDIT/SERVER=NEW_LOG
```

where *filespec* is a logical name that points to a node-specific file; for example, SYS\$SPECIFIC:[SYSMGR]SECURITY. System security audit log files on other nodes are unaffected.

Getting Information About the System

20.8 Monitoring Operating System Performance

20.8 Monitoring Operating System Performance

The Monitor utility (MONITOR) is a system management tool that you can use to obtain information about operating system performance. Various MONITOR qualifiers collect system performance data from the running system or play back data recorded previously in a recording file. When you play back data, you can display it, summarize it, and even rerecord it to reduce the amount of data in the recording file.

Following an explanation of the Monitor utility are sections that tell how to perform these tasks:

Task	Section
Invoking the Monitor utility	Section 20.8.2
Using live display monitoring	Section 20.8.3
Using live recording monitoring	Section 20.8.4
Using concurrent display and recording monitoring	Section 20.8.5
Using playback monitoring	Section 20.8.6
Using remote playback monitoring	Section 20.8.7
Rerecording monitoring	Section 20.8.8
Running MONITOR continuously	Section 20.8.9
Using remote monitoring	Section 20.8.10

For additional information about interpreting the information the Monitor utility provides, refer to the *OpenVMS Performance Management*. For additional information about using the Monitor utility, refer to the *OpenVMS System Management Utilities Reference Manual*.

20.8.1 Understanding MONITOR

Using MONITOR, you can monitor classes of systemwide performance data (such as system I/O statistics, page management statistics, and time spent in each of the processor modes) at specifiable intervals, and produce several types of output. You can also develop a database of performance information for your system by running MONITOR continuously as a background process, as explained in Section 20.8.9.

20.8.1.1 MONITOR Classes

Each MONITOR class consists of data items that, taken together, provide a statistical measure of a particular system performance category. The data items defined for individual classes are listed in the description of the MONITOR command in the *OpenVMS System Management Utilities Reference Manual*.

To monitor a particular class of information, specify a class name on the MONITOR command line. The information MONITOR displays depends on the type of class you select. Table 20–8 compares the two MONITOR class types.

Getting Information About the System

20.8 Monitoring Operating System Performance

Table 20–8 Types of MONITOR Classes

Type of class	Description
System	Provides statistics on resource use for the entire system
Component	Provides statistics on the contribution of individual components to the overall system or cluster

As an example of the distinction between types of MONITOR classes, the IO class includes a data item to measure all direct I/O operations for the entire system, and is therefore a system class. The DISK class measures direct I/O operations for individual disks, and is therefore a component class.

Table 20–9 describes each MONITOR class and indicates whether it is a system or component class.

Table 20–9 MONITOR Classes

Class	Type	Description
ALL_CLASSES	System or Component	Statistics for all classes
CLUSTER	System	Clusterwide performance statistics
DECNET	System	DECnet for OpenVMS statistics
DISK	Component	Disk I/O statistics
DLOCK	System	Distributed lock management statistics
FCP	System	File control primitive statistics
FILE_SYSTEM_CACHE	System	File system cache statistics
IO	System	System I/O statistics
LOCK	System	Lock management statistics
MODES	Component	Time spent in each of the processor modes
MSCP_SERVER	System	MSCP server statistics
PAGE	System	Page management statistics
PROCESSES	Component	Statistics on all processes
RMS	Component	Record Management Services statistics
SCS	Component	System Communications Services statistics
STATES	System	Number of processes in each of the scheduler states
SYSTEM	System	Summary of statistics from other classes
TRANSACTION	System	DECdtm services statistics
†VBS	System	Virtual balance slot statistics
VECTOR	System	Vector processor scheduled usage

†VAX specific

Getting Information About the System

20.8 Monitoring Operating System Performance

20.8.1.2 Display Data

Except in the PROCESSES class, all data item statistics are displayed as rates or levels:

- Rates are shown in number of occurrences per second.
- Levels are values that indicate the size of the monitored data item.

You can request any or all of four different statistics for each data item:

Statistic	Description
Current rate or level	Most recently collected value for the rate or level
Average rate or level	Measured from the beginning of the MONITOR request
Minimum rate or level	Measured from the beginning of the MONITOR request
Maximum rate or level	Measured from the beginning of the MONITOR request

For the DISK, MODES, SCS, and STATES classes, you can optionally express all statistics as percentages.

In the PROCESSES class, MONITOR displays descriptive information, level information, and counters that increase over time.

20.8.1.3 Output Types

MONITOR collects system performance data by class and produces three forms of optional output, depending on the qualifier you specify:

Qualifier	Description
/DISPLAY	Produces output in the form of ASCII screen images, which are written at a frequency governed by the /VIEWING_TIME qualifier.
/RECORD	Produces a binary recording file containing data collected for requested classes; one record for each class is written per interval.
/SUMMARY	Produces an ASCII file containing summary statistics for all requested classes over the duration of the MONITOR request.

If you specify /INPUT with any of these qualifiers, MONITOR collects performance data from one or more previously created recording files; otherwise, data is collected from counters and data structures on the running system.

You use the /BEGINNING and /ENDING qualifiers to specify, respectively, when you want a MONITOR request to begin and end.

Using the /DISPLAY Qualifier

Information collected by MONITOR is normally displayed as ASCII screen images. You can use the optional /DISPLAY qualifier to specify a disk file to contain the information. If you omit the file specification, output is directed to SYS\$OUTPUT.

Note

Be careful when you use the /DISPLAY qualifier. Because MONITOR enters display information into the file continuously, its size can grow very quickly.

Getting Information About the System

20.8 Monitoring Operating System Performance

Refer to the *OpenVMS System Management Utilities Reference Manual* for a discussion of the /DISPLAY qualifier.

Using the /RECORD Qualifier

When you use the /RECORD qualifier, all data pertaining to the class is recorded, even if you are concurrently displaying only a single statistic or a single item of a component statistics class. The file is created when a MONITOR request is initiated and closed when a request terminates. You can use the resulting file as a source file for later requests to format and display the data on a terminal, to create a summary file, or to create a new recording file with different characteristics.

20.8.2 Invoking MONITOR

To invoke the Monitor utility, enter the following DCL command:

```
$ MONITOR
```

MONITOR then displays the following prompt:

```
MONITOR>
```

In response to the prompt, you can enter any of the MONITOR commands, which are described in *OpenVMS System Management Utilities Reference Manual*. The most frequently used MONITOR command, however, specifies a class name.

Example

```
MONITOR> MONITOR PAGE
```

In this example, you specify the PAGE class name in the MONITOR command to monitor page management statistics.

You can also use the MONITOR command from DCL command level.

How to Override or Terminate a MONITOR Request

Generally, each MONITOR request runs until the time specified or implied by the /ENDING qualifier. However, to override or terminate a MONITOR request, you can press one of the following combinations of keys:

Keys	Description
Ctrl/W	Temporarily overrides a /VIEWING_TIME value and generates a new display immediately following the current one. This feature is useful when a broadcast message overwrites the MONITOR display area. You can also use Ctrl/W in conjunction with a large /VIEWING_TIME value to generate display events on demand.
Ctrl/C	Terminates the current request without exiting from the utility. You can then initiate a new request or enter any MONITOR command at the MONITOR> prompt.
Ctrl/Z	Terminates the current request and exits from MONITOR.

20.8.3 Using Live Display Monitoring

Use the live display monitoring mode of operation when you want to examine the activity of a running system, either on a routine basis or as part of an installation checkout, tuning, or troubleshooting exercise. The system does not keep a historical record of output. The following examples show how to use the live display monitoring mode.

Getting Information About the System

20.8 Monitoring Operating System Performance

Examples

1.

```
$ MONITOR PROCESSES/TOPCPU
```

The command displays a bar graph showing the eight processes that were the top consumers of CPU time during the period between displays. It also displays the amount of CPU time each process used.

The command might produce a display similar to the following one:

```

OpenVMS Monitor Utility
TOP CPU TIME PROCESSES
on node BOMBAY
20-JAN-2000 10:06:49

          0          25          50          75          100
+ - - - - + - - - - + - - - - + - - - - +
07E00181 CAFARET 100 *****
|
|
|
|
|
|
|
+ - - - - + - - - - + - - - - + - - - - +

```

This example shows that user CAFARET is using 100 percent of the CPU time available. To display more information about the computer resources a user is using, use a command similar to the following one:

```
$ SHOW PROCESS/CONTINUOUS/ID=07E00181
```

For this example, the most useful information in the resulting display is the name of the image at the end of the display; for example:

```

.
.
.
$1SDUA1:[SYS1D.SYSCOMMON.][SYSEXE]RODAN.EXE

```

This example indicates that CAFARET is running RODAN.EXE, which might be new software that is unintentionally running in a loop. This situation would occur if CAFARET were a privileged user running a process at a higher priority than other users.

2.

```
$ MONITOR/DISPLAY=PROCESSES.LOG PROCESSES
```

You can route MONITOR display output to any supported terminal device or to a disk file. This command writes MONITOR's display process statistics to the file PROCESSES.LOG. You can then print this file on a hardcopy device.

Caution

Because data is continuously added to the display file, be careful that the file does not grow too large.

Getting Information About the System

20.8 Monitoring Operating System Performance

3.

```
$ MY_CLASSES := -  
$ "DECNET+FCP+IO+LOCK+MODES+PAGE+PROCESSES+STATES"  
§ MONITOR/NODE=(CURLEY,LARRY)/INTERVAL=20/VIEWING_TIME=8 'MY_CLASSES'
```

You might find it convenient to establish DCL symbols for frequently used combinations of class names, as in this example. The MONITOR command collects selected classes of data for OpenVMS Cluster nodes CURLEY and LARRY every 20 seconds. Every 8 seconds, the system displays the most recently collected data for one of the classes. MONITOR predetermines the ordering of the classes for display.

20.8.4 Using Live Recording Monitoring

Use live recording to capture MONITOR data for future use. Possible uses include the following ones:

- Installation checkout, tuning, troubleshooting; that is, all the uses that are listed for live display monitoring.

Choose recording over display when you want to capture more classes than you can reasonably watch at a terminal, when a terminal is not available, or when you want to gather data about the system but cannot spend time at the terminal until later.

- Routine performance data gathering for long-term analysis.

You can record MONITOR data on a routine basis and summarize it to gather data about system resource use over long periods of time.

Caution

Because data is continuously added to the recording file, be careful that the file does not grow too large.

The following example shows how to use the live recording mode of operation.

Example

```
$ MONITOR/NODE=(LARRY,MOE)/NODISPLAY/RECORD MODES+STATES
```

The command in this example records data on the time spent in each processor mode and on the number of processes in each scheduler state for nodes LARRY and MOE. The command does not display this information.

20.8.5 Using Concurrent Display and Recording Monitoring

Use the concurrent display and recording mode of operation when you want to both retain performance data and watch as it is being collected. Because MONITOR allows shared read access to the recording file, a separate display process can play back the recording file as it is being written by another process.

The following examples show how to use the concurrent display and recording mode of operation. The first example both collects and records data in the same command. The second and third examples show how you can perform concurrent recording and display using two separate processes: the process in the second example performs recording; the process in the third example plays back the file to obtain a summary.

Getting Information About the System

20.8 Monitoring Operating System Performance

Examples

1.

```
$ MONITOR/RECORD FCP/AVERAGE,FILE_SYSTEM_CACHE/MINIMUM
```

This command collects and records file system data and file system cache data every 3 seconds. It also displays, in bar graph form, average FCP statistics and minimum FILE_SYSTEM_CACHE statistics. The display alternates between the two graphs every 3 seconds. You can obtain current statistics in a subsequent playback request.

2.

```
$ MONITOR/RECORD=SYS$MANAGER:ARCHIVE.DAT -  
_ $ /INTERVAL=300/NODISPLAY ALL_CLASSES
```

This command archives data for all classes once every 5 minutes. You might find it convenient to execute a similar command in a batch job, taking care to monitor disk space usage.

3.

```
$ MONITOR/INPUT=SYS$MANAGER:ARCHIVE.DAT: -  
_ $ /NODISPLAY/SUMMARY/BEGINNING="-1" PAGE,IO
```

The command in this example produces a summary of page and I/O activity that occurred during the previous hour, perhaps as part of an investigation of a reported performance problem. Note that because the recording process executes an OpenVMS RMS flush operation every 5 minutes, up to 5 minutes of the most recently collected data is not available to the display process.

You can specify the time between flush operations explicitly with the /FLUSH_INTERVAL qualifier. Note also that the display process must have read access to the recording file.

20.8.6 Using Playback Monitoring

Use playback of a recording file to obtain terminal output and summary reports of all collected data or a subset of it. You can make a subset of data according to class, node, or time segment. For example, if you collect several classes of data for an entire day, you can examine or summarize the data on one or more classes during any time period in that day.

You can also display or summarize data with a different interval than the one at which it was recorded. You control the actual amount of time between displays of screen images with the /VIEWING_TIME qualifier. The following examples show how to use the playback mode of operation.

Examples

1.

```
$ MONITOR/RECORD/INTERVAL=5 IO  
.  
.  
.  
$ MONITOR/INPUT IO
```

The commands in this example produce system I/O statistics. The first command gathers and displays data every 5 seconds, beginning when you enter the command and ending when you press Ctrl/Z. In addition, the first command records binary data in the default output file MONITOR.DAT.

Getting Information About the System

20.8 Monitoring Operating System Performance

The second command plays back the I/O statistics display, using the data in MONITOR.DAT for input. The default viewing time for the playback is 3 seconds, but each screen display represents 5 seconds of monitored I/O statistics.

2.

```
$ MONITOR/RECORD/NODISPLAY -
_ $ /BEGINNING=08:00:00 -
_ $ /ENDING=16:00:00 -
_ $ /INTERVAL=120 DISK
.
.
.
$ MONITOR/INPUT/DISPLAY=HOURLY.LOG/INTERVAL=3600 DISK
```

The sequence of commands in this example illustrates data recording with a relatively small interval and data playback with a relatively large interval. This is useful for producing average, minimum, and maximum statistics that cover a wide range of time, but have greater precision than if they had been gathered using the larger interval.

The first command records data on I/O operations for all disks on the system for the indicated 8-hour period, using an interval of 2 minutes. The second command plays the data back with an interval of 1 hour, storing display output in the file HOURLY.LOG. You can then type or print this file to show the cumulative average disk use at each hour throughout the 8-hour period.

Note

The current statistic in HOURLY.LOG shows the current data in terms of the original collection interval of 120 seconds, not the new collection interval of 3600 seconds.

3.

```
$ MONITOR/INPUT/NODISPLAY/SUMMARY=DAILY.LOG DISK
```

The command in this example uses the recording file created in the previous example to produce a one-page summary report file showing statistics for the indicated 8-hour period. The summary report has the same format as a screen display. For example:

```
OpenVMS Monitor Utility
DISK I/O STATISTICS
on node TLC
SUMMARY
From: 25-JAN-2000 08:00:00
To: 25-JAN-2000 16:00:00
```

I/O Operation Rate		CUR	AVE	MIN	MAX
DSA0:	SYSTEM 0	0.53	1.50	0.40	3.88
DSA1:	SYSTEM_1	0.00	0.39	0.00	8.38
DSA4:	WORK 0	0.00	0.11	0.00	1.29
DSA5:	WORK_1	0.03	0.87	0.00	5.95
DSA6:	WORK_2	0.03	0.25	0.00	2.69
DSA7:	WORK_3	0.04	0.97	0.00	20.33
DSA17:	TOM_DISK	0.00	0.04	0.00	0.80
DSA23:	MKC	0.00	0.00	0.00	0.13
\$4\$DUA0:	(RABBIT) SYSTEM 0	0.20	0.65	0.17	1.97
\$4\$DUA2:	(RABBIT) SYSTEM_0	0.20	0.65	0.17	1.97
\$4\$DUA3:	(RABBIT) SYSTEM_1	0.00	0.14	0.00	2.49

Getting Information About the System

20.8 Monitoring Operating System Performance

PLAYBACK

SUMMARIZING

20.8.7 Using Remote Playback Monitoring

If suitably privileged, you can collect MONITOR data from any system to which your system has a DECnet connection. You can then display the data live on your local system. To do so, follow these steps:

1. In the default DECnet directory on each remote system, create a file named MONITOR.COM, similar to the following example:

```
$ !
$ ! * Enable MONITOR remote playback *
$ !
$ MONITOR /NODISPLAY/RECORD=SYS$NET ALL_CLASSES
```

2. On your local system, define a logical name for the remote system from which you want to collect data. Use the following syntax:

```
DEFINE remotenodename_mon node::task=monitor
```

You might want to define, in a login command procedure, a series of logical names for all the systems you want to access.

3. To display the remote MONITOR data as it is being collected, enter a command using the following syntax:

```
MONITOR/INPUT=remotenodename_mon classnames
```

You can also place MONITOR.COM files in directories other than the default DECnet directory and use access control strings or proxy accounts to invoke these command files remotely.

When you invoke MONITOR on your local system, a process is created on the remote system that executes the MONITOR.COM command file. The remote system therefore experiences some associated CPU and DECnet overhead. You can regulate the overhead in the MONITOR.COM file by using the /INTERVAL qualifier and the list of class names.

Section 20.8.10 describes remote monitoring in a mixed-version cluster system.

20.8.8 Rerecording Monitoring

Rerecording is a combination of playback and recording. You can use it for data reduction of recording files. When you play back an existing recording file, all MONITOR options are available to you; thus, you can choose to record a subset of the recorded classes and a subset of the recorded time segment and a larger interval value.

All these techniques produce a new, smaller recording file at the expense of some of the recorded data. A larger interval value reduces the volume of the collected data, so displays and summary output produced from the newer recorded file will be less precise. Note that average rate values are not affected in this case, but average level values are less precise (since the sample size is reduced), as are maximum and minimum values. The following example shows how to use the rerecording mode of operation:

Getting Information About the System

20.8 Monitoring Operating System Performance

Example

```
$ SUBMIT MONREC.COM
```

MONREC.COM contains the following commands:

```
$ MONITOR/NODISPLAY/RECORD/INTERVAL=60 /BEGINNING=8:00/ENDING=16:00 DECNET,LOCK  
$ MONITOR/INPUT/NODISPLAY/RECORD DECNET
```

The first command runs in a batch job, recording DECnet and lock management data once every minute between the hours of 8 A.M. and 4 P.M.. The second command, which is issued after the first command completes, rerecords the data by creating a new version of the MONITOR.DAT file, containing only the DECnet data.

20.8.9 Running MONITOR Continuously

You can develop a database of performance information for your system by running MONITOR continuously as a background process. This section contains examples of procedures that you, as cluster manager, might use to create multifile clusterwide summaries.

You can adapt the command procedures to suit conditions at your site. Note that you must define the logical names SYS\$MONITOR and MON\$ARCHIVE in SYSTARTUP.COM before executing any of the command files.

The directory with the logical name SYS\$EXAMPLES includes three command procedures that you can use to establish the database. Instructions for installing and running the procedures are in the comments at the beginning of each procedure. Table 20–10 contains a brief summary of these procedures.

Table 20–10 MONITOR Command Procedures

Procedure	Description
MONITOR.COM	Creates a summary file from the recording file of the previous boot, and then begins recording for this boot. The recording interval is 10 minutes.
MONSUM.COM	Generates two clusterwide multifile summary reports that are mailed to the system manager: one report is for the previous 24 hours, and the other is for the previous day's prime-time period (9 A.M. to 6 P.M.). The procedure resubmits itself to run each day at midnight.
SUBMON.COM	Starts MONITOR.COM as a detached process. Invoke SUBMON.COM from the site-specific startup command procedure.

While MONITOR records data continuously, a summary report can cover any finite time segment. The MONSUM.COM command procedure, which is executed every midnight, produces and mails the two multifile summary reports described in Table 20–10. Because these reports are not saved as files, to keep them, you must either extract them from your mail file or alter the MONSUM.COM command procedure to save them.

20.8.9.1 Using the MONITOR.COM Procedure

The procedure in Example 20–7 archives the recording file and summary file from the previous boot and initiates continuous recording for the current boot. (Note that this procedure does not purge recording files.)

Getting Information About the System

20.8 Monitoring Operating System Performance

Example 20-7 MONITOR.COM Procedure

```
$ SET VERIFY
$ !
$ ! MONITOR.COM
$ !
$ ! This command file is to be placed in a cluster-accessible directory
$ ! called SYS$MONITOR and submitted at system startup time as a detached
$ ! process via SUBMON.COM. For each node, MONITOR.COM creates, in
$ ! SYS$MONITOR, a MONITOR recording file that is updated throughout the
$ ! life of the boot. It also creates, in MON$ARCHIVE, a summary file from
$ ! the recording file of the previous boot, along with a copy of that
$ ! recording file. Include logical name definitions for both cluster-
$ ! accessible directories, SYS$MONITOR and MON$ARCHIVE, in SYSTARTUP.COM.
$ !
$ SET DEF SYS$MONITOR
$ SET NOON
$ PURGE MONITOR.LOG/KEEP:2
$ !
$ ! Compute executing node name and recording and summary file names
$ ! (incorporating node name and date).
$ !
$ NODE = F$GETSYI("NODENAME")
$ SEP = ""
$ IF NODE .NES. "" THEN SEP = " "
$ DAY = F$EXTRACT(0,2,F$TIME())
$ IF F$EXTRACT(0,1,DAY) .EQS. " " THEN DAY = F$EXTRACT(1,1,DAY)
$ MONTH = F$EXTRACT(3,3,F$TIME())
$ ARCHFILNAM = "MON$ARCHIVE:"+NODE+SEP+"MON"+DAY+MONTH
$ RECFIL = NODE+SEP+"MON.DAT"
$ SUMFIL = ARCHFILNAM+".SUM"

$ !
$ ! Check for existence of recording file from previous boot and skip
$ ! summary if not present.
$ !
$ OPEN/READ/ERROR=NORECFIL RECORDING 'RECFIL'
$ CLOSE RECORDING
$ !
$ !
$ ! Generate summary file from previous boot.
$ !
$ MONITOR /INPUT='RECFIL' /NODISPLAY /SUMMARY='SUMFIL' -
$ ALL_CLASSES+MODE/ALL+STATES/ALL+SCS/ITEM=ALL+SYSTEM/ALL+DISK/ITEM=ALL
$ !
$ !
$ ! Compute subject string and mail summary file to cluster manager.
$ !
$ !
$ A=""
$ B=" MONITOR Summary "
$ SUB = A+NODE+B+F$TIME()+A
$ MAIL/SUBJECT='SUB' 'SUMFIL' CLUSTER_MANAGER
$ !
$ !
$ ! Archive recording file and delete it from SYS$MONITOR.
$ !
$ COPY 'RECFIL' 'ARCHFILNAM'.DAT
$ DELETE 'RECFIL';*
$ !
$ NORECFIL:
$ SET PROCESS/PRIORITY=15
$ !
```

(continued on next page)

Getting Information About the System

20.8 Monitoring Operating System Performance

Example 20–7 (Cont.) MONITOR.COM Procedure

```
$ !
$ ! Begin recording for this boot. The specified /INTERVAL value is
$ ! adequate for long-term summaries; you might need a smaller value
$ ! to get reasonable "semi-live" playback summaries (at the expense
$ ! of more disk space for the recording file).
$ !
$ MONITOR /INTERVAL=300 /NODISPLAY /RECORD='RECFIL' ALL_CLASSES
$ !
$ !
$ ! End of MONITOR.COM
$ !
```

20.8.9.2 Using the SUBMON.COM Procedure

The procedure in Example 20–8 submits MONITOR.COM as a detached process from SYSTARTUP.COM to initiate continuous recording for the current boot.

Example 20–8 SUBMON.COM Procedure

```
$ SET VERIFY
$ !
$ ! SUBMON.COM
$ !
$ ! This command file is to be placed in a cluster-accessible directory
$ ! called SYS$MONITOR. At system startup time, for each node, it is
$ ! executed by SYSTARTUP.COM, following logical name definitions for
$ ! the cluster-accessible directories SYS$MONITOR and MON$ARCHIVE.
$ !
$ !
$ ! Submit detached MONITOR process to do continuous recording.
$ !
$ !
$ !
$ RUN  SYS$SYSTEM:LOGINOUT.EXE -
        /UIC=[1,4] -
        /INPUT=SYS$MONITOR:MONITOR.COM -
        /OUTPUT=SYS$MONITOR:MONITOR.LOG -
        /ERROR=SYS$MONITOR:MONITOR.LOG -
        /PROCESS_NAME="Monitor" -
        /WORKING_SET=512 -
        /MAXIMUM_WORKING_SET=512 -
        /EXTENT=512/NOSWAPPING
$ !
$ !
$ ! End of SUBMON.COM
$ !
```

20.8.9.3 Using the MONSUM.COM Procedure

The procedure in Example 20–9 produces daily and prime-time clusterwide summaries.

Getting Information About the System

20.8 Monitoring Operating System Performance

Example 20-9 MONSUM.COM Procedure

```
$ SET VERIFY
$ !
$ ! MONSUM.COM
$ !
$ ! This command file is to be placed in a cluster-accessible directory
$ ! called SYS$MONITOR and executed at the convenience of the cluster
$ ! manager. The file generates both 24-hour and "prime time" cluster
$ ! summaries and resubmits itself to run each day at midnight.
$ !
$ SET DEF SYS$MONITOR
$ SET NOON
$ !
$ ! Compute file specification for MONSUM.COM and resubmit the file.
$ !
$ FILE = F$ENVIRONMENT("PROCEDURE")
$ FILE = F$PARSE(FILE,,,"DEVICE")+F$PARSE(FILE,,,"DIRECTORY")+F$PARSE(FILE,,,"NAME")
$ SUBMIT 'FILE' /AFTER=TOMORROW /NOPRINT
$ !
$ ! Generate 24-hour cluster summary.
$ !
$ !
$ MONITOR/INPUT=(SYS$MONITOR:*MON*.DAT;* ,MON$ARCHIVE:*MON*.DAT;*) -
  /NODISPLAY/SUMMARY=MONSUM.SUM -
  ALL CLASSES+DISK/ITEM=ALL+SCS/ITEM=ALL-
  /BEGIN="YESTERDAY+0:0:0.00" /END="TODAY+0:0:0.00" /BY_NODE
$ !
$ !
$ ! Mail 24-hour summary file to cluster manager and delete the file from
$ ! SYS$MONITOR.
$ !
$ !
$ MAIL/SUBJECT="Daily Monitor Clusterwide Summary" MONSUM.SUM CLUSTER_MANAGER
$ DELETE MONSUM.SUM;*
$ !
$ ! Generate prime-time cluster summary.
$ !
$ !
$ MONITOR/INPUT=(SYS$MONITOR:*MON*.DAT;* ,MON$ARCHIVE:*MON*.DAT;*) -
  /NODISPLAY/SUMMARY=MONSUM.SUM -
  ALL CLASSES+DISK/ITEM=ALL+SCS/ITEM=ALL-
  /BEGIN="YESTERDAY+9:0:0.00" /END="YESTERDAY+18:0:0.00" /BY_NODE
$ !
$ !
$ ! Mail prime-time summary file to cluster manager and delete the file
$ ! from SYS$MONITOR.
$ !
$ !
$ MAIL/SUBJECT="Prime-Time Monitor Clusterwide Summary" MONSUM.SUM CLUSTER_MANAGER
$ DELETE MONSUM.SUM;*
$ !
$ ! End of MONSUM.COM
$ !
```

Note that Mail commands in this procedure send files to user `CLUSTER_MANAGER`. Replace `CLUSTER_MANAGER` with the appropriate user name or logical name for your site.

Because summary data might be extensive, Compaq recommends that you print out summary files.

20.8.10 Using Remote Monitoring

MONITOR is capable of using both TCP/IP and DECnet as a transport mechanism. Beginning with OpenVMS Version 7.0, to use TCP/IP, you must start the TCP/IP server by issuing the following command inside SYS\$STARTUP:SYSTARTUP_VMS.COM:

```
$ @SYS$STARTUP:VPM$STARTUP.COM
```

DECnet continues to work as in the past: a network object is created at the time of the request.

Remote Monitoring in a Mixed-Version OpenVMS Cluster System

You can monitor any node in an OpenVMS Cluster system either by issuing the MONITOR CLUSTER command or by adding the /NODE qualifier to any interactive MONITOR request.

Remote monitoring in an OpenVMS Cluster system might not be compatible, however, between nodes that are running different versions of OpenVMS. Table 20–11 shows the compatibility of versions for remote monitoring.

Table 20–11 Remote Monitoring Compatibility in an OpenVMS Cluster System

Versions	OpenVMS Alpha and VAX Version 6.n or 7.n	OpenVMS Alpha Version 1.5 and VAX Version 5.n
OpenVMS Alpha and VAX Version 6.n or 7.n	Yes	No
OpenVMS Alpha Version 1.5 and VAX Version 5.n	No	Yes

If you attempt to monitor a remote node that is incompatible, the system displays the following message:

```
%MONITOR-E-SRVMISMATCH, MONITOR server on remote node is an incompatible version
```

If this is the case, contact your Compaq support representative for a remedial kit that corrects this problem. Before you install the remedial kit, you can still use MONITOR to obtain data about the remote node. To do this, record the data on the remote node and then run the MONITOR playback feature to examine the data on the local node.

Another difference exists when you monitor remote nodes in an OpenVMS Cluster system. Beginning with OpenVMS Version 6.2, the limit on the number of disks that can be monitored was raised from 799 to 909 for record output and from 799 to 1817 for display and summary outputs. If you monitor a remote node running OpenVMS Version 6.2 or later from a system running a version earlier than OpenVMS Version 6.2, the old limit of 799 applies.

For more information about MONITOR, refer to the *OpenVMS System Management Utilities Reference Manual*.

Tracking Resource Use

This chapter describes how to find out how your system resources have been used. You can use this information to:

- Charge users for the resources they have used. You can produce reports of the resources used by individual users.
- Plan your future equipment requirements. You can monitor changing patterns of resource use and predict future demands.
- Troubleshoot the system. You can check the final exit status of processes.
- Improve system performance. You can find out the load that individual images and processes place on your system.
- Detect security breaches. You can identify unusual patterns of resource use.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Determining which resources are being tracked	Section 21.2
Controlling which resources are tracked	Section 21.3
Starting up a new accounting file	Section 21.4
Moving the accounting file	Section 21.5
Producing reports of resource use	Section 21.6
Setting up accounting groups	Section 21.7
Monitoring disk space	Section 21.8

This chapter explains the following concept:

Concept	Section
Accounting files	Section 21.1

21.1 Understanding Accounting Files

The system gathers information about resource use. For example, the information can include the resources such as CPU time used by each print job. The system stores this information in **accounting files**.

The resources tracked by default depend on the model of computer you use. However, you can control which resources are tracked. If you do not want to track resource use, you can stop the accounting file tracking resource use altogether. (See Section 21.3.)

Tracking Resource Use

21.1 Understanding Accounting Files

Each node in an OpenVMS Cluster has its own accounting file, known as its **current accounting file**. By default, this file is SYS\$MANAGER:ACCOUNTNG.DAT, but you can control which file is used (see Section 21.5).

The information in the accounting files is in binary. You cannot display it with the TYPE command. To display the information, use the Accounting utility (ACCOUNTING). (See Section 21.6.)

21.2 Determining Which Resources Are Being Tracked

To determine which resources are currently being tracked, use the SHOW ACCOUNTING command:

```
$ SHOW ACCOUNTING
```

This command produces a screen display (see the example) that contains keywords in the following two categories:

- Keywords that show which types of resource are being tracked:

Keyword	Type of Resource
IMAGE	Resources used by an image
LOGIN_FAILURE	Resources used by an unsuccessful attempt to log in
MESSAGE	Unformatted resource record written to the accounting file by a call to the \$SNDJBC system service
PRINT	Resources used by a print job
PROCESS	Resources used by a process

- Keywords that show which types of process are being tracked. When the resources for processes or images are tracked, these keywords show the process type:

Keyword	Type of Process
BATCH	Batch process
DETACHED	Detached process
INTERACTIVE	Interactive process
NETWORK	Network process
SUBPROCESS	Subprocess (the parent process can be a batch, detached, interactive, or network process)

Example

```
$ SHOW ACCOUNTING
```

Accounting is currently enabled to log the following activities:

```
PROCESS      any process termination
IMAGE        image execution
INTERACTIVE  interactive job termination
LOGIN_FAILURE login failures
NETWORK      network job termination
PRINT        all print jobs
```

The keywords in this example show that the local node is tracking the resources used by each:

Tracking Resource Use

21.2 Determining Which Resources Are Being Tracked

- Interactive and network process
- Image running in an interactive or network process
- Login failure
- Print job

21.3 Controlling Which Resources Are Tracked

You can control which resources the system tracks. To save disk space, you can stop the system tracking resources you are not interested in.

How to Perform This Task

1. Use the SET ACCOUNTING command with the /ENABLE and /DISABLE qualifiers in the following format to control which resources are tracked:

```
SET ACCOUNTING/DISABLE[(keyword[,...])]/ENABLE[(keyword[,...])]
```

The keywords are the same as those explained in Section 21.2.

2. To make this change permanent, edit the SET ACCOUNTING command in the SYS\$MANAGER:SYSTART_VMS.COM startup file.

Example

This example prevents the tracking of all resources except those used by interactive and batch processes:

```
$ SET ACCOUNTING/DISABLE/ENABLE=(PROCESS,INTERACTIVE,BATCH)
```

The /DISABLE qualifier is not followed by a keyword. Therefore, the qualifier disables the tracking of all resources. The /ENABLE qualifier then enables the tracking of the resources used by interactive and batch processes.

21.4 Starting Up a New Accounting File

To start up a new current accounting file, use the following command:

```
$ SET ACCOUNTING/NEW_FILE
```

This closes the current accounting file and opens a new version of it.

If the system encounters an error when trying to write information to the current accounting file, it automatically closes the file and opens a new version of it.

Example

This example closes the current accounting file, opens a new version of it, and changes the name of the old file to WEEK_24_RESOURCES.DAT. You can retain this file as a record of the resources used in that week.

```
$ SET ACCOUNTING/NEW_FILE  
$ RENAME SYS$MANAGER:ACCOUNTNG.DAT;-1 WEEK_24_RESOURCES.DAT
```

21.5 Moving the Accounting File

When you first install your system, the current accounting file is SYS\$MANAGER:ACCOUNTNG.DAT.

This file can become quite large. Moving it from your system disk can improve system performance.

Tracking Resource Use

21.5 Moving the Accounting File

How to Perform This Task

1. Define the logical name ACCOUNTNG in your system logical name table to point to the file you want to use. For example:

```
$ DEFINE ACCOUNTNG MYDISK:[MYDIR]MYFILE.DAT/SYSTEM
```

Give the full file specification, including the device and directory.

Note

Two nodes cannot log information in the same accounting file. If you define ACCOUNTNG on two nodes to point to the same file, each node will open and use its own version of the file.

2. To make the change permanent, add this definition to the file SYS\$MANAGER:SYLOGICALS.COM.
3. Use the SET ACCOUNTING command with the /NEW_FILE qualifier to create and use the new file:

```
$ SET ACCOUNTING/NEW_FILE
```

Example

This example changes the current accounting file to [MYDIR]MYDISK:MYFILE.DAT.

```
$ DEFINE ACCOUNTNG MYDISK:[MYDIR]MYFILE.DAT/SYSTEM
$ SET ACCOUNTING/NEW_FILE
```

21.6 Producing Reports of Resource Use

The three types of reports are:

Type of Report	Qualifier
Brief	/BRIEF (the default)
Full	/FULL
Summary	/SUMMARY

To produce a report, use the ACCOUNTING command with the appropriate qualifier in the following format:

```
ACCOUNTING [filespec[...]/qualifier[...]]
```

This runs the Accounting utility. The *filespec* parameter lists the accounting files you want to process. If you omit it, the Accounting utility processes the default current accounting file, SYS\$MANAGER:ACCOUNTNG.DAT.

By default, the Accounting utility processes all the records in the accounting files you specify. You can use selection qualifiers to specify which records you want to process.

By default, brief and full reports present the records in the order in which they were logged in the accounting file. When you produce brief and full reports, you can use the /SORT qualifier to specify another order.

Tracking Resource Use 21.6 Producing Reports of Resource Use

Example

This example produces a brief report of the information in the file that the logical name ACCOUNTNG points to. The /TYPE qualifier selects records for print jobs only. The /SORT qualifier displays them in reverse alphabetical order of user name.

```
$ ACCOUNTING ACCOUNTNG/TYPE=PRINT/SORT=-USER
```

Date / Time	Type	Subtype	Username	ID	Source	Status
13-APR-2000 13:36:04	PRINT		SYSTEM	20A00442		00000001
13-APR-2000 12:42:37	PRINT		JONES	20A00443		00000001
13-APR-2000 14:43:56	PRINT		FISH	20A00456		00000001
14-APR-2000 19:39:01	PRINT		FISH	20A00265		00000001
14-APR-2000 20:09:03	PRINT		EDWARDS	20A00127		00000001
14-APR-2000 20:34:45	PRINT		DARNELL	20A00121		00000001
14-APR-2000 11:23:34	PRINT		CLARK	20A0032E		00040001
14-APR-2000 16:43:16	PRINT		BIRD	20A00070		00040001
14-APR-2000 09:30:21	PRINT		ANDERS	20A00530		00040001

21.7 Setting Up Accounting Groups

Users are already organized into UIC security groups. For accounting purposes, security groups are often inappropriate. You can put users into accounting groups with the Authorize utility using the /ACCOUNT qualifier. In this way, each user is in an accounting group and a security group.

Using the Accounting utility, you can:

- Summarize the resources used by all the users in a particular accounting or security group. To do this, use the ACCOUNT or UIC keyword with the /SUMMARY qualifier.
- Select records for all the users in a particular accounting or security group. To do this, use the /ACCOUNT or /UIC qualifier.

How to Perform This Task

1. Plan your accounting groups. Decide which users you want in each accounting group, and choose names for the groups.

The name of an accounting group can be a maximum of eight characters.

2. Change the account field values in the UAF. Use the Authorize utility's MODIFY command in the following format to change the value in the account field to the name of the user's accounting group:

```
MODIFY username/ACCOUNT=account-name
```

where:

username is the name of the user

account-name is the name of the accounting group that you want that user to be in

The next time your users log in, they will be in their new accounting groups, and their resource use will be tagged with the appropriate accounting group names.

Tracking Resource Use

21.7 Setting Up Accounting Groups

Example

This example modifies the accounting group name to SALES_W8 for the username FORD:

```
$ RUN SYS$SYSTEM:AUTHORIZE
UAF> MODIFY FORD/ACCOUNT=SALES_W8
UAF> EXIT
```

21.8 Monitoring Disk Space

To find out how much disk space a user is using, use `SYSMAN` or, if you have not enabled disk quotas, the `DIRECTORY` command.

How to Perform This Task

Use either of the following methods:

- Use the `SYSMAN` command `DISKQUOTA SHOW` in the following format:

```
DISKQUOTA SHOW owner [/DEVICE=device-spec]
```

This shows the number of blocks used by all the files that are owned by the specified user on the specified disk.

- Use the `DIRECTORY` command with the `/SIZE` and `/GRAND_TOTAL` qualifiers in the following format:

```
DIRECTORY [filespec[...]]/SIZE=ALLOCATION/GRAND_TOTAL
```

This shows the number of blocks used by all the files in the specified file location.

Note that the `DIRECTORY` command does not include the blocks used by file headers or the user's root directory.

Examples

1. This example uses `SYSMAN` to find out the number of blocks used by all the files that are owned by each user.

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> DISKQUOTA SHOW *

%SYSMAN-I-QUOTA, disk quota statistics on device SYS$SYSTEM:MYDISK
Node UNION
      UIC                Usage      Permanent Quota  Overdraft Limit
[0,0]                0          1000             100
[DOC,EDWARDS]       115354     150000           5000
[DOC,FISH]          177988     250000           5000
[DOC,SMITH]         140051     175000           5000
[DOC,JONES]         263056     300000           5000
```

2. This example uses the `DIRECTORY` command to show the number of blocks allocated by all the files in and under `MYDISK:[PARSONS]`.

```
$ DIRECTORY MYDISK:[PARSONS...]/SIZE=ALLOCATION/GRAND_TOTAL
Grand total of 28 directories, 2546 files, 113565 blocks.
```

OpenVMS Cluster Considerations

This chapter describes concepts related to the OpenVMS Cluster environment; it also tells how the Show Cluster utility (SHOW CLUSTER) can display information about a cluster and how the System Management utility (SYSMAN) can help you manage an OpenVMS Cluster environment.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Beginning to use SHOW CLUSTER commands	Section 22.3.2
Adding information to a report	Section 22.3.3
Controlling the display of data	Section 22.3.4
Formatting the display of data	Section 22.3.5
Creating a startup initialization file	Section 22.3.6
Using command procedures containing SHOW CLUSTER commands	Section 22.3.7
Using SYSMAN to manage security	Section 22.5
Using the SYSMAN DO command to manage an OpenVMS Cluster	Section 22.6

This chapter explains the following concepts:

Concept	Section
OpenVMS Cluster systems	Section 22.1
Setting up an OpenVMS Cluster environment	Section 22.1.1
Clusterwide system management	Section 22.1.2
The Show Cluster utility (SHOW CLUSTER)	Section 22.3.1
SYSMAN and OpenVMS Cluster management	Section 22.4

22.1 Understanding OpenVMS Cluster Systems

An **OpenVMS Cluster system** is a loosely coupled configuration of two or more computers and storage subsystems, including at least one Alpha computer. An OpenVMS Cluster system appears as a single system to the user even though it shares some or all of the system resources. When a group of computers shares resources clusterwide, the storage and computing resources of all of the computers are combined, which can increase the processing capability, communications, and availability of your computing system.

OpenVMS Cluster Considerations

22.1 Understanding OpenVMS Cluster Systems

A **shared resource** is a resource (such as a disk) that can be accessed and used by any node in an OpenVMS Cluster system. Data files, application programs, and printers are just a few items that can be accessed by users on a cluster with shared resources, without regard to the particular node on which the files or program or printer might physically reside.

When disks are set up as shared resources in an OpenVMS Cluster environment, users have the same environment (password, privileges, access to default login disks, and so on) regardless of the node that is used for logging in. You can realize a more efficient use of mass storage with shared disks, because the information about any device can be used by more than one node—the information does not have to be rewritten in many places. You can use the OpenVMS MSCP, which is the mass storage control protocol, or TMSCP, which is the tape mass storage control protocol, server software to make tapes accessible to nodes that are not directly connected to the storage devices.

You can also set up print and batch queues as shared resources. In an OpenVMS Cluster configuration with shared print and batch queues, a single queue database manages the queues for all nodes. The queue database makes the queues available from any node. For example, suppose your cluster configuration has fully shared resources and includes nodes ALBANY, BASEL, and CAIRO. A user logged in to node ALBANY can send a file that physically resides on node BASEL to a printer that is physically connected to node CAIRO, and the user never has to specify (or even know) the nodes for either the file or the printer.

Planning an OpenVMS Cluster System

A number of types of OpenVMS Cluster configurations are possible. Refer to *Guidelines for OpenVMS Cluster Configurations* and the OpenVMS Cluster Software Product Description (SPD) for complete information about supported devices and configurations.

The following sections briefly describe OpenVMS Cluster systems. For complete information about setting up and using an OpenVMS Cluster environment, refer to *OpenVMS Cluster Systems*.

22.1.1 Setting Up an OpenVMS Cluster Environment

Once you have planned your configuration, installed the necessary hardware, and checked hardware devices for proper operation, you can set up an OpenVMS Cluster system using various system software facilities. Setup procedures to build your cluster follow.

Procedure	For More Information
Installing or upgrading the operating system on the first OpenVMS Cluster computer	Installation and operations guide for your computer
Installing required software licenses	<i>OpenVMS License Management Utility Manual</i>
Configuring and starting the DECnet for OpenVMS network	<i>DECnet for OpenVMS Networking Manual</i>
Configuring and starting TCP/IP Services	<i>Compaq TCP/IP Services for OpenVMS Installation and Configuration</i>

OpenVMS Cluster Considerations

22.1 Understanding OpenVMS Cluster Systems

Procedure	For More Information
Preparing files that define the cluster operating environment and that control disk and queue operations	<i>OpenVMS Cluster Systems</i>
Adding computers to the cluster	<i>OpenVMS Cluster Systems</i>

Depending on various factors, the order in which these operations are performed can vary from site to site, as well as from cluster to cluster at the same site.

22.1.2 Clusterwide System Management

Once any system is installed, you must decide how to manage users and resources for maximum productivity and efficiency while maintaining the necessary security. OpenVMS Cluster systems provide the flexibility to distribute users and resources to suit the needs of the environment. OpenVMS Cluster system resources can also be easily redistributed as needs change. Even with the vast number of resources available, you can manage the cluster configuration as a single system.

You have several tools and products to help you manage your cluster as a unified entity.

OpenVMS Cluster Tools

The following utilities are provided with the operating system:

Utility	Description
Compaq DECams	Collects and analyzes data from multiple nodes simultaneously, directing all output to a centralized DECwindows display. (Refer to Section 22.2 and the <i>DECams User's Guide</i> .)
Monitor utility (MONITOR)	Provides basic performance data. (See Section 20.8.)
Show Cluster utility (SHOW CLUSTER)	Monitors activity in an OpenVMS Cluster configuration, and then collects and sends information about that activity to a terminal or other output device. (Described in Section 22.3.)
System Management utility (SYSMAN)	Allows you to send common control commands across all, or a subset of, the nodes in the cluster. (Described in Section 22.6.)

System Management Applications

The following products are *not* provided with the operating system:

Product	Description
POLYCENTER solutions	A comprehensive set of operations management products and services to help you manage complex distributed environments. However, the POLYCENTER Software Installation utility is described in this manual in Section 3.7.
†Storage Library System (SLS) for VAX ‡Storage Library System (SLS) for Alpha	A set of software tools that enables tape, cartridge tape, and optical disks.
OpenVMS Cluster Console System (VCS)	Designed to consolidate the console management of the OpenVMS Cluster system at a single console terminal.

†VAX specific

‡Alpha specific

OpenVMS Cluster Considerations

22.1 Understanding OpenVMS Cluster Systems

You can find additional information about these system management tools in the appropriate product documentation.

22.2 Using Compaq DECams to Analyze Data

Compaq DECams is a real-time monitoring, diagnostic, and correction tool that assists system managers to improve OpenVMS system and OpenVMS Cluster availability. DECams can help system programmers and analysts target a specific node or process for detailed analysis, and can help system operators and service technicians resolve hardware and software problems.

DECams simultaneously collects and analyzes system data and process data from multiple nodes and displays the output on a DECwindows Motif display. Based on the collected data, DECams analyzes, detects, and proposes actions to correct resource and denial issues in real-time.

For more information, refer to the *DECams User's Guide*.

22.3 Using SHOW CLUSTER

The Show Cluster utility (SHOW CLUSTER) monitors nodes in an OpenVMS Cluster system. You can use the utility to display information about cluster activity and performance.

The following sections describe the Show Cluster utility and explain how to perform these tasks:

Task	Section
Begin to use SHOW CLUSTER commands	Section 22.3.2
Add information to a report	Section 22.3.3
Control the display of data	Section 22.3.4
Format the display of data	Section 22.3.5
Create a startup initialization file	Section 22.3.6
Use command procedures containing SHOW CLUSTER commands	Section 22.3.7

22.3.1 Understanding SHOW CLUSTER

You can display SHOW CLUSTER information on your terminal screen or send it to a device or a file. You can use SHOW CLUSTER interactively, with command procedures, or with an initialization file in which you define default settings. Because this utility is installed with the CMKRNL privilege, SHOW CLUSTER requires no special privilege.

SHOW CLUSTER information includes approximately 100 fields of data. You can customize the appearance of SHOW CLUSTER reports or define reports for access to often-needed data.

SHOW CLUSTER reports are organized by classes and fields:

Unit of Organization	Description
Class	Group of related fields of information. You can use class names to selectively add or remove an entire class from a report. Each class displays certain fields by default. Some classes have additional fields that you can add or remove using the field name.
Field	Column of data in a report. Use a unique field name to refer to each field of data. You can use the field name to selectively add or remove a field from reports. For the names and descriptions of all of the fields in each class, see the ADD (Field) command in the <i>OpenVMS System Management Utilities Reference Manual</i> .

You can add fields or classes to the default SHOW CLUSTER report. If you add a field or class to a report in a continuous display, SHOW CLUSTER automatically adds the new data to the display.

Figure 22–1 shows a sample default SHOW CLUSTER report. The default report has two classes of information: SYSTEMS and MEMBERS. Below each class name are columns of fields that are associated with each class of information.

Figure 22–1 SHOW CLUSTER Default Display

View of Cluster from system ID 65536 node: CLUB 31-DEC-1997 14:00:00

SYSTEMS		MEMBERS
NODE	SOFTWARE	STATUS
CLUB	VMS V7.2	MEMBER
HSJ400	HSJ V25J	
HSC900	HSC V860	
CHIP	VMS V7.1	MEMBER
DISK3	RFX V256	
DISK1	RFX V256	
SPREE	VMS V6.2	MEMBER
SPRITZ	VMS V7.1	MEMBER

ZK-8998A-GE

Table 22–1 briefly describes the fields shown in Figure 22–1.

Table 22–1 Fields in Default SHOW CLUSTER Report

Field Name	Description
NODE	Node name of the remote system. Normally, the cluster manager sets the node name using the system parameter SCSNODE. The node name should be the same as the DECnet for OpenVMS node name.
SOFTWARE	Name and version of the operating system currently running on the remote system.
STATUS	Status of the node in the cluster. (MEMBER indicates that the system is participating in the cluster.)

Over time, you can determine the most valuable classes and fields of data for your SHOW CLUSTER reports; you can then create a startup initialization file that establishes your default report formats. You can also build command procedures to use while running SHOW CLUSTER interactively. In this way,

OpenVMS Cluster Considerations

22.3 Using SHOW CLUSTER

you can quickly reformat the report to show the data that is relevant for your installation. Startup initialization files and command procedures are explained later in this chapter.

Because SHOW CLUSTER information includes many fields of data, the report can quickly extend beyond screen limits. Therefore, SHOW CLUSTER provides mechanisms to help you control the display of data, including the following mechanisms:

- 38 SHOW CLUSTER commands
- A default keypad, which you can redefine

These mechanisms are described in detail in the *OpenVMS System Management Utilities Reference Manual*.

22.3.2 Beginning to Use SHOW CLUSTER Commands

To use the Show Cluster utility, enter the SHOW CLUSTER command. If you specify the command without any qualifiers, however, SHOW CLUSTER simply displays a default report like that shown in Figure 22–1 and then displays the DCL prompt.

In a continuous display, on the other hand, you can enter SHOW CLUSTER commands to control report output. You can, for example, add classes or fields to, or remove classes or fields from, reports. To invoke a continuous display, in which you can enter SHOW CLUSTER commands, use the /CONTINUOUS qualifier on the SHOW CLUSTER command. (SHOW CLUSTER command qualifiers are described in Section 22.3.2.3.)

How to Perform This Task

To invoke a continuous display of default SHOW CLUSTER report information, enter the following command:

```
$ SHOW CLUSTER/CONTINUOUS
```

SHOW CLUSTER then displays a default report. By default, SHOW CLUSTER updates the display every 15 seconds, with the changed data displayed in reverse video. After the default report, SHOW CLUSTER displays the following prompt:

```
Command>
```

(If the report extends below the limit of your terminal screen and you do not see the Command> prompt, you can press Return to display the prompt.)

The following sections contain instructions for performing beginning SHOW CLUSTER tasks:

Task	Section
Viewing information that is off the screen	Section 22.3.2.1
Exiting from a continuous display	Section 22.3.2.2
Using SHOW CLUSTER qualifiers	Section 22.3.2.3

22.3.2.1 Viewing Information That Is Off the Screen

The PAN command allows you to view the entire display by shifting your view of the display by column (horizontally) or by line (vertically).

Note

Report headings also move out of view as the reports in the display are panned beyond the limits of the screen. The SCROLL command, which is explained in Section 22.3.5.4, preserves the headings as you scroll information. To use the SCROLL command, you must take the additional step of selecting a report if you have more than one report on the screen.

How to Perform This Task

To pan the display, perform one of the following actions:

- Enter PAN commands at the command prompt; for example:

```
Command> PAN DOWN 10
```

The command in this example moves the display down 10 lines.

- Define the arrow keys as PAN commands:

```
Command> SET FUNCTION PAN
```

This command redefines the arrow keys as follows:

Arrow Key	Redefinition
UP ARROW KEY	PAN UP 1
DOWN ARROW KEY	PAN DOWN 1
→	PAN RIGHT 1
←	PAN LEFT 1

You can then use the arrow keys to move up, down, right, and left in the display.

Refer to the SET FUNCTION and PAN commands in the *OpenVMS System Management Utilities Reference Manual* for details.

Resetting Arrow Keys

By default, the SHOW CLUSTER arrow keys are set to the EDIT function. This means that, at the command prompt, you can perform command line editing that is similar to DCL line-mode editing. For example, the left arrow key moves the cursor to the left, and the up arrow key recalls the previous command. Refer to the *OpenVMS User's Manual* for information about DCL line-mode editing.

When you use the SET FUNCTION command, you reset the function keys. After that, the arrow keys are redefined and DCL line-mode editing is disabled.

To reset the arrow keys, enter the following command:

```
Command> SET FUNCTION EDIT
```

22.3.2.2 Exiting from a Continuous Display

To exit from a continuous display, perform one of the following actions:

- To return to the DCL prompt, perform one of the following actions:
 - Enter EXIT after the Command> prompt.
 - Press Ctrl/Z.
 - Press Ctrl/Y.

OpenVMS Cluster Considerations

22.3 Using SHOW CLUSTER

- To exit without erasing the screen, press Ctrl/C.

22.3.2.3 Using SHOW CLUSTER Qualifiers

Table 22–2 briefly describes the qualifiers you can use with the SHOW CLUSTER command. The *OpenVMS System Management Utilities Reference Manual* contains reference information about these SHOW CLUSTER qualifiers.

Table 22–2 SHOW CLUSTER Qualifiers

Qualifier	Function
/BEGINNING= <i>time</i>	Specifies the time that the SHOW CLUSTER session is to begin.
/CONTINUOUS	Controls whether SHOW CLUSTER runs as a continuously updating display.
/ENDING= <i>time</i>	Specifies the time that the SHOW CLUSTER session is to end.
/INTERVAL= <i>seconds</i>	Specifies the number of seconds that report information remains on the screen before it is updated.
/OUTPUT= <i>file-spec</i>	Directs the output from SHOW CLUSTER to the specified file instead of to the current SYS\$OUTPUT device.

Example

In a continuous display, SHOW CLUSTER updates the display every 15 seconds by default. You can change this interval by using the /INTERVAL qualifier.

```
$ SHOW CLUSTER/CONTINUOUS/INTERVAL=5
```

In this example, SHOW CLUSTER updates reports every 5 seconds, displaying changed data in reverse video.

22.3.3 Adding Information to a Report

When you use the SHOW CLUSTER command, the resulting report is only part of the total information available. As shown in Figure 22–1, the default classes displayed are MEMBERS and SYSTEMS. Table 22–3 briefly describes all the classes you can display in SHOW CLUSTER reports. Refer to the *OpenVMS System Management Utilities Reference Manual* for details about these classes.

Table 22–3 Classes of Information Available in SHOW CLUSTER Reports

Classes	Information Displayed
CIRCUITS	Describes virtual circuits on OpenVMS Cluster systems.
CLUSTER	Shows general information about the OpenVMS Cluster system, such as the time it was formed, the last time a system joined or left, and the cluster quorum.
CONNECTIONS	Describes the connections established over a virtual circuit in the OpenVMS Cluster system
COUNTERS	Shows counts of the total accumulated traffic over a connection for the life of the connection.
CREDITS	Shows send and receive credit counts for connections in the OpenVMS Cluster system.

(continued on next page)

Table 22–3 (Cont.) Classes of Information Available in SHOW CLUSTER Reports

Classes	Information Displayed
ERRORS	Displays a count of the errors on each port, along with information about the feasibility of reinitializing a port.
LOCAL_PORTS	Displays information about the local system interface to the OpenVMS Cluster system, such as the name, number, and status of each port, and the number of entries in the queues associated with each port.
MEMBERS	Describes systems actively participating in the OpenVMS Cluster system.
SYSTEMS	Describes all OpenVMS Cluster systems. It shows node name, identification number, hardware type, and software version.

Example

The following example shows how to add the CLUSTER class to a SHOW CLUSTER display:

```
Command> ADD CLUSTER
```

Figure 22–2 shows the display that results from entering the ADD CLUSTER command. CLUSTER class is displayed below the default SHOW CLUSTER display.

Figure 22–2 SHOW CLUSTER Display with CLUSTER Report

View of Cluster from system ID 65536 node: CLUB 31-DEC-1997 14:00:00

SYSTEMS		MEMBERS
NODE	SOFTWARE	STATUS
CLUB	VMS V7.2	MEMBER
HSJ400	HSJ V25J	
HSC900	HSC V860	MEMBER
CHIP	VMS V7.1	
DISK3	RFX V256	
DISK1	RFX V256	MEMBER
SPREE	VMS V6.2	
SPRITZ	VMS V7.1	MEMBER

CLUSTER						
CL_EXP	CL_QUORUM	CL_VOTES	QF_VOTE	CL_MEMBERS	FORMED	LAST_TRANSITION
3	2	3	NO	4	15-JUN-1997	10-DEC-1997

ZK-8999A-GE

For descriptions of the fields in the CLUSTER class, refer to the SHOW CLUSTER section of the *OpenVMS System Management Utilities Reference Manual*.

OpenVMS Cluster Considerations

22.3 Using SHOW CLUSTER

22.3.4 Controlling the Display of Data

Using SHOW CLUSTER commands, you can remove fields or classes from a display, remove broadcast messages from the screen, and refresh the screen display at any time. The following sections explain how to perform these operations.

22.3.4.1 Entering Commands to Display Data

SHOW CLUSTER allows you to customize the display of data during a continuous session by entering various commands. The *OpenVMS System Management Utilities Reference Manual* describes SHOW CLUSTER commands in detail.

Updating of the continuous display stops as soon as you enter input from the terminal keyboard. When you press the Return key after entering a command, updating of the display resumes until you enter another command.

By default, updating takes place at 15-second intervals. If you do not enter a new command within 15 seconds, the command prompt disappears, and two more lines of data take its place.

22.3.4.2 Removing Broadcast Messages

When you receive a system broadcast message during a continuous SHOW CLUSTER session, the message appears at the bottom of the screen. A multiline message fills as many lines of the screen as it needs.

How to Perform This Task

The last broadcast message you receive remains on the screen until you acknowledge it by entering input from the terminal in one of the following ways:

- Press the Return key.
- Refresh the screen by pressing Ctrl/W.
- Enter a command.

If you receive more than one broadcast message, SHOW CLUSTER waits until the next update interval to display the next message.

SHOW CLUSTER also displays error messages at the bottom of the screen. For an explanation of the error messages, refer to the *OpenVMS System Messages: Companion Guide for Help Message Users*.

22.3.4.3 Refreshing the Screen

Ordinarily, a continuous display is updated or refreshed according to the default or specified interval time. SHOW CLUSTER scans the software databases, extracts and stores data for each field, displays any new or changed data, and updates the time. On Compaq and Compaq-compatible terminals, reverse video highlights any changed data.

How to Perform This Task

You can refresh the screen at any time by one of the following methods:

- Modify the format of the display with the ADD, REMOVE, INITIALIZE, or SET command.
- Use the REFRESH command.
- Press Ctrl/W.

22.3.5 Formatting the Display of Data

Because SHOW CLUSTER allows you to include additional fields and classes, you can produce reports that overflow the physical limits of the terminal screen. However, you can use a number of methods to modify the display to meet your needs:

Formatting Method	For More Information
Remove data from reports	Section 22.3.5.1
Modify field and screen size	Section 22.3.5.2
Move a report	Section 22.3.5.3
Scroll a report	Section 22.3.5.4

22.3.5.1 Removing Information from a Report

You can remove certain fields or classes to reduce the width of a report to fit the limits of your screen. Also, certain fields or classes might not be important for your particular needs. You can also remove particular types of data to reduce the length of the report.

How to Perform This Task

Use the REMOVE command to remove fields or entire classes. To remove a field or class, use the appropriate qualifier with the REMOVE command. refer to the REMOVE commands in the SHOW CLUSTER section of the *OpenVMS System Management Utilities Reference Manual* for appropriate class names and qualifiers.

Examples

1.

```
Command> REMOVE SOFTWARE
```

The command in this example removes the SOFTWARE field from the SHOW CLUSTER report shown in Figure 22–1.

Refer to the ADD (Field) command description in the *OpenVMS System Management Utilities Reference Manual* for a list of valid field names.

2.

```
Command> REMOVE MEMBERS
```

The command in this example removes the MEMBERS class from the SHOW CLUSTER report shown in Figure 22–1.

22.3.5.2 Modifying Field and Screen Size

To make a report fit the physical limits of the screen, you can change the width of certain fields in the report. For example, if SHOW CLUSTER provides a field width that can contain any possible value and the values your cluster generates do not require that much space, you can adjust the field width with the SET (Field) command.

SHOW CLUSTER also allows you to adjust the size of the terminal screen. If the terminal is Compaq-compatible and supports a wide report, you can set the screen to a width of up to 511 columns by specifying an appropriate value to the SET SCREEN command.

OpenVMS Cluster Considerations

22.3 Using SHOW CLUSTER

Examples

1.

```
Command> SET TRANSITION_TYPE/WIDTH=10
```

The command in this example sets the width of the TRANSITION_TYPE field to 10, which removes the time of day from the field but leaves the date.

2.

```
Command> SET SCREEN=132
```

The command in this example sets the screen width to 132.

Refer to the *OpenVMS System Management Utilities Reference Manual* for more details about using the SET (Field) and SET SCREEN commands.

22.3.5.3 Moving a Report

By default, SHOW CLUSTER operates with AUTO_POSITIONING ON. This means that the utility automatically arranges the reports to take best advantage of the available display space. However, you can position reports manually with the MOVE command, which implicitly sets AUTO_POSITIONING to OFF.

If you have multiple reports in your display, you must first select the report to be repositioned. You use the SELECT *window-name* command to specify the report name; for example:

- SCS (the default report, which usually includes fields in the SYSTEMS and MEMBERS classes)
- CLUSTER
- LOCAL_PORTS

Note

To select any report except the default SCS report, you must first add the class to the display if it is not already displayed; for example:

```
Command> ADD LOCAL_PORTS
```

As an alternative, you can repeatedly press the Select function key or the period key on the keypad to cycle from one report to the next. The selected report appears highlighted.

How to Perform This Task

To move a report, perform either of the following actions:

- Enter MOVE commands at the command prompt.
- Use the arrow keys that you define as MOVE commands.

```
Command> SET FUNCTION MOVE
```

This command redefines the arrow keys as follows:

Arrow Key	Redefinition
UP ARROW KEY	MOVE UP 1

Arrow Key	Redefinition
DOWN ARROW KEY	MOVE DOWN 1
→	MOVE RIGHT 1
←	MOVE LEFT 1

When you enter a MOVE command, the display changes position by column (horizontally) or by line (vertically). For example, entering the command MOVE LEFT 5 moves the display 5 columns to the left. An empty frame appears around the new position of the report.

When you are satisfied with the position of the report, enter the DESELECT command, which moves the report to the new position. Entering another SELECT command before the previous MOVE operation has been deselected also moves the report to its new position.

Example

```
Command> SELECT CLUSTER
Command> MOVE RIGHT 10
Command> DESELECT
```

The following lists explains the commands in the example:

1. The SELECT command selects the CLUSTER report (which is then highlighted).
2. The MOVE command positions the report frame 10 spaces to the right.
3. The DESELECT command terminates the MOVE operation and displays the contents of the report.

For more information, refer to the SELECT, SET FUNCTION, and DESELECT commands in the *OpenVMS System Management Utilities Reference Manual*.

To reset the arrow keys, enter the following command:

```
Command> SET FUNCTION EDIT
```

22.3.5.4 Scrolling a Report

The SCROLL command provides a means of quickly scanning through a report without losing column headings. Scrolling scans a display by field (horizontally) and by line (vertically). The report headings remain stationary when you scroll vertically.

When the display has more than one report, you must first select a report by entering the SELECT command. The selected report is highlighted.

How to Perform This Task

To scroll a display, perform either of the following actions:

- Enter SCROLL commands at the command prompt.
- Use the arrow keys that you define as SCROLL commands.

```
Command> SET FUNCTION SCROLL
```

OpenVMS Cluster Considerations

22.3 Using SHOW CLUSTER

This command redefines the arrow keys as follows:

Arrow Key	Redefinition
UP ARROW KEY	SCROLL UP 1
DOWN ARROW KEY	SCROLL DOWN 1
→	SCROLL RIGHT 1
←	SCROLL LEFT 1

Example

```
Command> SELECT SCS
Command> SET FUNCTION SCROLL
```

The commands in this example first select the SCS report (which is then highlighted), and then set the arrow keys to scroll functions. Refer to the SET FUNCTION and SCROLL commands in the *OpenVMS System Management Utilities Reference Manual* for more information.

To reset the arrow keys, enter the following command:

```
Command> SET FUNCTION EDIT
```

22.3.6 Creating a Startup Initialization File

To customize the SHOW CLUSTER display, you can create a startup initialization file, which the utility executes when you enter it. SHOW CLUSTER takes the original default display, and adds or removes whatever classes or fields you specify. The resulting display becomes your default startup format. A startup initialization file resembles the following example:

```
!
!Startup Initialization File
!
!
INITIALIZE
REMOVE MEMBERS
ADD RP_REVISION,RP_TYPE,SYS_ID
SET SCREEN=132
```

This startup procedure deletes the MEMBERS class information from the default display. The procedure also adds the RP_REVISION and RP_TYPE fields from the CIRCUITS class and the SYS_ID field from the SYSTEMS class. The last line of the procedure sets the screen size to 132 columns.

How to Perform This Task

To create an initialization file, follow these steps:

1. Define the logical name SHOW_CLUSTER\$INIT as *device:[directory]SHCINI* before invoking SHOW CLUSTER.

For a startup file to execute before the display begins, you must assign the logical name SHOW_CLUSTER\$INIT to the initialization file; for example:

```
DEFINE SHOW_CLUSTER$INIT DEVA:[JONES]SHCINI
```

When invoked, SHOW CLUSTER searches for the file defined by SHOW_CLUSTER\$INIT. In this example, SHOW CLUSTER looks for DEVA:[JONES]SHCINI.INI when it starts up. If the initialization file is found, SHOW CLUSTER executes the procedure before beginning the display.

If you do not define SHOW_CLUSTER\$INIT or it does not include a directory specification, SHOW CLUSTER searches the current default directory for a file named SHOW_CLUSTER.INI.

2. Customize the display using SHOW CLUSTER commands during a continuous SHOW CLUSTER session.
3. Preserve the command sequence by entering the following command:

```
Command> SAVE SHOW_CLUSTER$INIT.INI
```

You must specify SHOW_CLUSTER\$INIT.INI, because the SAVE command creates a file with a file type of .COM by default. SHOW CLUSTER looks for an .INI file when it searches for a startup initialization file.

You can edit the file that the SAVE command creates to include comments or to improve its efficiency. For more information, refer to the SAVE command in the *OpenVMS System Management Utilities Reference Manual*.

Instead of having SHOW CLUSTER build an initialization file, you can build one yourself in the same way you build a command procedure. The next section provides guidelines for creating a command procedure.

22.3.7 Using Command Procedures Containing SHOW CLUSTER Commands

You can create command procedures that contain SHOW CLUSTER commands. Such files let you modify display characteristics without having to enter commands interactively. You can use command procedures during a continuous SHOW CLUSTER session to perform a series of commands, for example, to customize the output of the display.

The following list contains guidelines for writing command procedures that contain SHOW CLUSTER commands:

- Use any valid SHOW CLUSTER commands.
- Nest command procedures up to 16 levels deep.
- Include the SHOW CLUSTER command INITIALIZE as the first command in the file. The INITIALIZE command ensures that the report is in a known state before any commands are executed to modify it.

Notes

Do not include an EXIT command at the end of the command procedure. The EXIT command terminates SHOW CLUSTER and erases the SHOW CLUSTER display before you can see it.

Also, do not run SHOW CLUSTER command procedures from a batch job.

The following command procedure customizes a report display:

```
!  
! Include only the node field from the default display; show votes  
! and quorum for each node and for the cluster as a whole.  
!  
INITIALIZE  
REMOVE SOFTWARE,STATUS  
ADD VOTES,QUORUM,CL_VOTES,CL_QUORUM
```

OpenVMS Cluster Considerations

22.3 Using SHOW CLUSTER

This command procedure removes the SOFTWARE and STATUS fields from the report and adds fields that provide information about the cluster quorum and votes.

To execute a command procedure during a continuous SHOW CLUSTER session, specify the execute procedure (@) command, along with the file name of the command procedure. The default file type for command procedure files is .COM.

Example

The following command executes a command procedure named SYSMOD.COM:

```
Command> @SYSMOD
```

In this example, the default file type .COM is assumed because the file type is omitted.

For more information about creating command procedures, refer to the SAVE command in the *OpenVMS System Management Utilities Reference Manual*.

22.4 Understanding SYSMAN and OpenVMS Cluster Management

The System Management utility (SYSMAN) provides two kinds of support for OpenVMS Cluster management:

- Cluster-specific commands, CONFIGURATION SET and CONFIGURATION SHOW, that you can use to manage security data and system time in a cluster
- Access to DCL-level commands with the DO command, which gives you the ability to apply a single DCL command across an entire cluster, rather than having to enter the command on each node

Each SYSMAN command requires a specific level of privilege. For more information about each command, refer to the *OpenVMS System Management Utilities Reference Manual*.

22.5 Using SYSMAN to Manage Security

You can manage security data for an OpenVMS Cluster system with SYSMAN CONFIGURATION commands. Table 22–4 summarizes these CONFIGURATION commands and their functions.

Table 22–4 SYSMAN CONFIGURATION Commands

Command	Function
CONFIGURATION SET CLUSTER_AUTHORIZATION	Modifies the group number and password in a local area cluster
CONFIGURATION SHOW CLUSTER_AUTHORIZATION	Displays the group number and multicast address of a local area cluster

22.5.1 Modifying the Group Number and Password

The group number identifies the group of nodes in the cluster, and the associated Ethernet address is used to send messages to all nodes in the cluster. The OpenVMS Cluster password protects the integrity of the cluster membership.

Using the CONFIGURATION SET CLUSTER_AUTHORIZATION command modifies the group number and password, as recorded in SYS\$SYSTEM:CLUSTER_AUTHORIZE.DAT. Normally, you do not need to alter records in the CLUSTER_AUTHORIZE.DAT file.

OpenVMS Cluster Considerations 22.5 Using SYSMAN to Manage Security

If your configuration has multiple system disks, SYSMAN automatically updates each copy of CLUSTER_AUTHORIZE.DAT, provided that you have defined the environment as a cluster with the SET ENVIRONMENT/CLUSTER command.

Caution

If you change either the group number or password, you must reboot the entire cluster.

You cannot display the cluster password for security reasons, but you can display the group number and group multicast address with the CONFIGURATION SHOW CLUSTER_AUTHORIZATION command.

Examples

1. The following command example sets the environment to a specific cluster, sets privilege to SYSPRV, and modifies the cluster password:

```
SYSMAN> SET ENVIRONMENT/CLUSTER/NODE=NODE21
SYSMAN> SET PROFILE/PRIVILEGE=SYSPRV
SYSMAN> CONFIGURATION SET CLUSTER AUTHORIZATION/PASSWORD=GILLIAN
%SYSMAN-I-CAFOLDGROUP, existing group will not be changed
%SYSMAN-I-GRPNOCHG, Group number not changed
SYSMAN-I-CAFREBOOT, cluster authorization file updated.
The entire cluster should be rebooted.
```

2. The following command example displays the group number and multicast address for NODE21. Because the group number and password on other nodes in the cluster are identical, no further information is displayed.

```
SYSMAN> CONFIGURATION SHOW CLUSTER AUTHORIZATION
Node NODE21: Cluster group number 65240
Multicast address: AB-00-04-01-F2-FF
```

22.6 Using the SYSMAN Command DO to Manage an OpenVMS Cluster

The SYSMAN command DO enables you to execute a DCL command or command procedure on all nodes in the current environment. This is convenient when you are performing routine system management tasks on nodes in the OpenVMS Cluster system, such as:

- Installing images
- Starting up software
- Checking devices
- Checking memory

Each DO command executes as an independent process, so there is no process context retained between DO commands. For this reason, you must express all DCL commands in a single command string, and you cannot run a program that expects input.

In a cluster environment, SYSMAN executes the commands sequentially on all nodes in the cluster. Each command executes completely before SYSMAN sends it to the next node in the environment. Any node that is unable to execute the command returns an error message. SYSMAN displays an error message if the timeout period expires before the node responds.

OpenVMS Cluster Considerations

22.6 Using the SYSMAN Command DO to Manage an OpenVMS Cluster

In a dual-architecture heterogeneous OpenVMS Cluster running both OpenVMS VAX and OpenVMS Alpha, some uses of the DO command may require special handling. For example, if you are installing images that are named differently in each architecture, you can still use the DO command if you create logical name tables for VAX and for Alpha nodes. See the example sequence that follows this description for an example.

Some DCL commands, such as MOUNT/CLUSTER or SET QUORUM/CLUSTER, operate clusterwide by design. Similarly, operations on clusterwide logical names and tables operate clusterwide by design. It is best to avoid using these kinds of commands with the DO command in SYSMAN when the environment is set to cluster. As alternatives, you could leave SYSMAN temporarily with the SPAWN command and execute these commands in DCL, or you could define the environment to be a single node within the cluster.

Examples

1. The following example installs an image on a cluster. First, it adds CMKRNL and SYSPRV privileges to the current privileges because they are required by INSTALL and AUTHORIZE. The DO INSTALL command installs the file STATSHR. The DO MCR AUTHORIZE command sets up an account for user Jones, specifying a password and a default device and directory.

```
SYSMAN> SET PROFILE/PRIVILEGES=(CMKRNL,SYSPRV)/DEFAULT=SYS$SYSTEM
SYSMAN> DO INSTALL ADD/OPEN/SHARED WRKD$:[MAIN]STATSHR
SYSMAN> DO MCR AUTHORIZE ADD JONES/PASSWORD=COLUMBINE -
_SYSMAN> /DEVICE=WORK1/DIRECTORY=[JONES]
```

2. The following example sets the environment to cluster and starts up a software product called XYZ on each node in the cluster:

```
SYSMAN>SET ENVIRONMENT/CLUSTER
%SYSMAN-I-ENV, Current command environment:
      Clusterwide on local cluster
      Username SMITH will be used on nonlocal nodes
SYSMAN> DO @SYS$STARTUP:XYZ_STARTUP
```

3. The following example shows how you can define logical names for VAX and Alpha nodes in a dual-architecture heterogeneous cluster, so that you can use the DO command to install architecture-specific images.

```
$ CREATE/NAME TABLE/PARENT=LNMS$SYSTEM DIRECTORY SYSMAN$NODE TABLE
$ DEFINE/TABLE=SYSMAN$NODE TABLE ALPHA_NODES NODE21,NODE22,NODE23
$ DEFINE/TABLE=SYSMAN$NODE_TABLE VAX_NODES NODE24,NODE25,NODE26
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/NODE=ALPHA_NODES
%SYSMAN-I-ENV, current command environment:
      Individual nodes: NODE21,NODE22,NODE23
      Username BOUCHARD will be used on nonlocal nodes

SYSMAN> DO INSTALL REPLACE SYS$LIBRARY:DCLTABLES.EXE
%SYSMAN-I-OUTPUT, command execution on node NODE21
%SYSMAN-I-OUTPUT, command execution on node NODE22
%SYSMAN-I-OUTPUT, command execution on node NODE23
SYSMAN> DO INSTALL REPLACE SYS$SYSTEM: DEC FORTRAN.EXE
%SYSMAN-I-OUTPUT, command execution on node NODE21
%SYSMAN-I-OUTPUT, command execution on node NODE22
%SYSMAN-I-OUTPUT, command execution on node NODE23

SYSMAN> SET ENVIRONMENT/NODE=VAX_NODES
%SYSMAN-I-ENV, current command environment:
      Individual nodes: NODE24,NODE25,NODE26
      Username BOUCHARD will be used on nonlocal nodes
```

OpenVMS Cluster Considerations

22.6 Using the SYSMAN Command DO to Manage an OpenVMS Cluster

```
SYSMAN> DO INSTALL REPLACE SYS$LIBRARY:DCLTABLES.EXE
%SYSMAN-I-OUTPUT, command execution on node NODE24
%SYSMAN-I-OUTPUT, command execution on node NODE25
%SYSMAN-I-OUTPUT, command execution on node NODE26
SYSMAN> DO INSTALL REPLACE SYS$SYSTEM:FORTRAN$MAIN.EXE
%SYSMAN-I-OUTPUT, command execution on node NODE24
%SYSMAN-I-OUTPUT, command execution on node NODE25
%SYSMAN-I-OUTPUT, command execution on node NODE26
```

4. The following example shows which files are open on DISK2. You might use this if you want to dismount DISK2 and need to see which users in the cluster have files open.

```
SYSMAN >SET ENVIRONMENT/CLUSTER
%SYSMAN-I-ENV, Current command environment:
    Clusterwide on local cluster
    Username SMITH will be used on nonlocal nodes
SYSMAN> DO SHOW DEVICE/FILES DISK2:

%SYSMAN-I-OUTPUT, command execution on node NODE21
Files accessed on device $1$DIA2: (DISK2, NODE22) on 14-MAY-2000 15:44:06.05
Process name      PID      File name
                00000000  [000000]INDEXF.SYS;1
%SYSMAN-I-OUTPUT, command execution on node NODE22
Files accessed on device $1$DIA2: (DISK2, NODE21) on 14-MAY-2000 15:44:26.93
Process name      PID      File name
                00000000  [000000]INDEXF.SYS;1
%SYSMAN-I-OUTPUT, command execution on node NODE23
Files accessed on device $1$DIA2: (NODE21, NODE22) on 14-MAY-2000 15:45:01.43
Process name      PID      File name
                00000000  [000000]INDEXF.SYS;1
%SYSMAN-I-OUTPUT, command execution on node NODE24
Files accessed on device $1$DIA2: (NODE22, NODE21) on 14-MAY-2000 15:44:31.30
Process name      PID      File name
                00000000  [000000]INDEXF.SYS;1
Susan Scott      21400059  [SCOTT]DECW$SM.LOG;228
_FTA7:          214000DD  [SCOTT]CARE_SDML.TPU$JOURNAL;1
%SYSMAN-I-OUTPUT, command execution on node NODE25
Files accessed on device $1$DIA2: (NODE21, NODE22) on 14-MAY-2000 15:44:35.50
Process name      PID      File name
                00000000  [000000]INDEXF.SYS;1
DECW$SESSION    226000E6  [SNOW]DECW$SM.LOG;6
_FTA17:        2260009C  [SNOW.MAIL]MAIL.MAI;1
_SNOW_1        2260012F  [SNOW.MAIL]MAIL.MAI;1
_SNOW_2        22600142  [SNOW.MAIL]MAIL.MAI;1
_SNOW_3        22600143  [SNOW.MAIL]MAIL.MAI;1
```

OpenVMS Cluster Considerations

22.6 Using the SYSMAN Command DO to Manage an OpenVMS Cluster

5. The following example shows how much memory is available on the nodes in a cluster. You might use this if you are installing software and want to know if each node has enough memory available.

```
SYSMAN > SET ENVIRONMENT/NODE=(NODE21,NODE22)
%SYSMAN-I-ENV, Current command environment:
      Clusterwide on local cluster
      Username SMITH      will be used on nonlocal nodes
SYSMAN> DO SHOW MEMORY
%SYSMAN-I-OUTPUT, command execution on node NODE21
      System Memory Resources on 14-MAY-2000 15:59:21.61
Physical Memory Usage (pages):      Total      Free      In Use      Modified
Main Memory (64.00Mb)              131072    63955     65201      1916
Slot Usage (slots):                Total      Free      Resident    Swapped
Process Entry Slots                 360       296       64         0
Balance Set Slots                   324       262       62         0
Fixed-Size Pool Areas (packets):    Total      Free      In Use      Size
Small Packet (SRP) List             10568     1703      8865       128
I/O Request Packet (IRP) List       3752      925       2827       176
Large Packet (LRP) List             157       28        129       1856
Dynamic Memory Usage (bytes):       Total      Free      In Use      Largest
Nonpaged Dynamic Memory             1300480   97120     1203360    60112
Paged Dynamic Memory                1524736  510496    1014240    505408
Paging File Usage (pages):          Free      Reservable Total
DISK$MTWAIN_SYS:[SYS0.SYSEXE]SWAPFILE.SYS
                                     10000     10000     10000
DISK$MTWAIN_SYS:[SYS0.SYSEXE]PAGEFILE.SYS
                                     60502     -52278    100000
Of the physical pages in use, 19018 pages are permanently allocated to VMS.

%SYSMAN-I-OUTPUT, command execution on node NODE22
      System Memory Resources on 14-MAY-2000 15:59:42.65
Physical Memory Usage (pages):      Total      Free      In Use      Modified
Main Memory (32.00Mb)              65536     44409     20461      666
Slot Usage (slots):                Total      Free      Resident    Swapped
Process Entry Slots                 240       216       24         0
Balance Set Slots                   212       190       22         0
Fixed-Size Pool Areas (packets):    Total      Free      In Use      Size
Small Packet (SRP) List             5080      2610     2470       128
I/O Request Packet (IRP) List       3101      1263     1838       176
Large Packet (LRP) List             87        60        27       1856
Dynamic Memory Usage (bytes):       Total      Free      In Use      Largest
Nonpaged Dynamic Memory             1165312  156256    1009056    114432
Paged Dynamic Memory                1068032  357424    710608    352368
Paging File Usage (pages):          Free      Reservable Total
DISK$MTWAIN_SYS:[SYS1.SYSEXE]SWAPFILE.SYS
                                     10000     10000     10000
DISK$MTWAIN_SYS:[SYS1.SYSEXE]PAGEFILE.SYS
                                     110591     68443    120000
Of the physical pages in use, 9056 pages are permanently allocated to VMS.
```

Network Considerations

This chapter discusses the following subjects:

- TCP/IP and DECnet networking software options for OpenVMS systems
- How to decide which networking software options are appropriate for your OpenVMS system
- How to prepare your system to join a network
- Where to obtain detailed information to help you install, configure, and manage the networking software you choose

The material provided in this chapter is intended as an introduction only. For complete planning, installation, configuration, use, and management information about the networking products, refer to the applicable product documentation.

Information Provided in This Chapter

This chapter explains the following concepts:

Concept	Section
Networking software options for OpenVMS systems	Section 23.1
Introduction to Compaq TCP/IP Services for OpenVMS software	Section 23.3
Introduction to DECnet-Plus for OpenVMS software	Section 23.8
Using DECnet over TCP/IP	Section 23.11

This chapter describes the following tasks:

Task	Section
Choosing the networking software appropriate for your system	Section 23.2
Preparing to join a TCP/IP network	Section 23.4
Installing and configuring TCP/IP Services	Section 23.5
Starting and stopping TCP/IP Services	Section 23.6
Preparing to join a DECnet-Plus network	Section 23.9
Installing and configuring DECnet-Plus	Section 23.10
Moving from DECnet Phase IV to DECnet-Plus	Section 23.12
Starting and stopping DECnet-Plus	Section 23.13
Running DECnet applications over an IP network backbone	Section 23.11

Network Considerations

23.1 OpenVMS Networking Software Options

23.1 OpenVMS Networking Software Options

The ability to connect your OpenVMS system with other systems and networks is a fundamental part of the OpenVMS operating system.

The following networking software options are available for OpenVMS systems:

- Compaq TCP/IP Services for OpenVMS
- TCP/IP implementations for OpenVMS from vendors other than Compaq
- Compaq DECnet-Plus for OpenVMS (Phase V)
- Compaq DECnet for OpenVMS (Phase IV)

You can install and use one of these networking options on an OpenVMS system, or you can install and use combinations of the options to accomplish the following:

- Interoperate with and share resources with other OpenVMS systems
- Interoperate with and share resources with systems running other operating systems, such as UNIX and Windows NT
- Use DECnet protocols and applications in combination with TCP/IP protocols and applications

23.2 Choosing Compaq Networking Software

Table 23–1 describes the Compaq layered networking software for OpenVMS systems. The software is available for you to install during the OpenVMS operating system installation or upgrade procedure. Or, you can install the software separately as layered products.

Table 23–1 Choosing Compaq Networking Software for OpenVMS Systems

Product	Description
Compaq TCP/IP Services for OpenVMS	<p>TCP/IP Services is the Compaq implementation of the industry-standard Transmission Control Protocol and Internet Protocol (TCP/IP) suite of protocols and Internet services for OpenVMS VAX and Alpha systems. If your OpenVMS system needs to communicate with heterogeneous networks (such as the Internet, UNIX systems, and Windows NT systems), you should choose Compaq TCP/IP Services for OpenVMS software. With TCP/IP Services, you can connect to remote hosts and access files, exchange messages, develop applications, monitor the network, and perform other important tasks.</p> <p>You can use TCP/IP Services on your system as your only networking software. OpenVMS does not require that you use DECnet software.</p> <p>TCP/IP protocols can coexist with DECnet-Plus protocols. You can install TCP/IP Services to serve as your network backbone and install DECnet-Plus as well, so that you can continue to use DECnet applications and functionality. For information about using DECnet protocols over an IP backbone, see Section 23.11.</p> <p>For an introduction to installing, using, and managing TCP/IP Services software, see Section 23.3.</p>

(continued on next page)

Network Considerations

23.2 Choosing Compaq Networking Software

Table 23–1 (Cont.) Choosing Compaq Networking Software for OpenVMS Systems

Product	Description
DECnet-Plus for OpenVMS (Phase V)	<p>DECnet-Plus is the Compaq Phase V implementation of the Digital Network Architecture (DNA). The software provides full backward compatibility with the older DECnet Phase IV product as well as the ability to run DECnet (NSP) and OSI over a DECnet, OSI, or TCP/IP network backbone. For information about using DECnet and OSI protocols over an IP backbone, see Section 23.11.</p> <p>If your OpenVMS system needs to communicate using DECnet (Phase IV and Phase V) or OSI applications and protocols, and if your system needs to coexist with TCP/IP or OSI protocols, or both, you should choose DECnet-Plus software.</p> <p>If you are gradually upgrading your network from DECnet Phase IV, the DECnet-Plus features help you by allowing you to use Phase IV functionality as well as the newer OSI features and advantages.</p> <p>For an introduction to installing, using, and managing DECnet-Plus software, see Section 23.8.</p> <p><u>Note:</u> Because DECnet-Plus contains both DNA Phase IV and Phase V protocols, you cannot run the separate DECnet Phase IV software product on the same system on which you install DECnet-Plus.</p> <p>For information about upgrading to DECnet-Plus, see the <i>DECnet-Plus Planning Guide</i>.</p>
DECnet for OpenVMS (Phase IV)	<p>DECnet Phase IV is the Compaq Phase IV implementation of the Digital Network Architecture (DNA). The product does not include the OSI protocols and the TCP/IP communications capability of the later DECnet-Plus (Phase V) product. You can run TCP/IP Services software on your system along with DECnet Phase IV, but you cannot run IP as the network backbone.</p> <p>If you intend to use your OpenVMS system to communicate in the traditional OpenVMS environment rather than a heterogeneous environment, you might want to choose DECnet Phase IV, the older DECnet product.</p> <p>This chapter does not contain summary product information about the older DECnet Phase IV product. For information about using DECnet Phase IV, refer to the <i>DECnet for OpenVMS Networking Manual</i> and the <i>DECnet for OpenVMS Network Management Utilities</i> manual (available on the OpenVMS Documentation CD-ROM).</p>

Table 23–2 lists possible networking software combinations for communication between OpenVMS systems.

Table 23–2 Compaq Networking Software Interoperability Options

If System A has...	And System B has...	Systems A and B can communicate using...
TCP/IP Services	TCP/IP Services	TCP/IP applications
DECnet Phase IV	DECnet Phase IV	DECnet applications
DECnet-Plus	DECnet-Plus	DECnet applications OSI applications
DECnet-Plus	DECnet Phase IV	DECnet applications

(continued on next page)

Network Considerations

23.2 Choosing Compaq Networking Software

Table 23–2 (Cont.) Compaq Networking Software Interoperability Options

If System A has...	And System B has...	Systems A and B can communicate using...
DECnet-Plus	OSI	OSI applications
TCP/IP Services and DECnet Phase IV	TCP/IP Services	TCP/IP applications
TCP/IP Services and DECnet Phase IV	DECnet Phase IV	DECnet applications
TCP/IP Services and DECnet-Plus	TCP/IP Services	TCP/IP applications
TCP/IP Services and DECnet-Plus	DECnet-Plus	DECnet applications OSI applications
TCP/IP Services and DECnet-Plus	TCP/IP Services and DECnet-Plus	OSI applications DECnet applications DECnet applications using DECnet over TCP/IP OSI applications using OSI over TCP/IP TCP/IP applications
TCP/IP Services and DECnet-Plus	OSI (supporting RFC 1006) and TCP/IP Services	OSI applications OSI over TCP/IP TCP/IP applications
TCP/IP Services and DECnet-Plus	OSI (not supporting RFC 1006) and TCP/IP Services	OSI applications TCP/IP applications

23.3 Understanding Compaq TCP/IP Services for OpenVMS

The Compaq TCP/IP Services for OpenVMS product is the OpenVMS implementation of the industry-standard TCP/IP suite of communications protocols as specified in Request for Comments (RFCs) used by the Internet Engineering Task Force (IETF).

With TCP/IP, heterogeneous networks can interconnect, making it possible for users to connect to remote hosts in many ways:

- Network file access. Users can access files on remote hosts.
- Electronic mail. Users can exchange messages between hosts.
- Application development. Application programmers can develop TCP/IP client/server applications for communication between local and remote hosts.
- Download and file transfer. Users can exchange files between hosts.
- User information. Users can access information about other users logged onto the local or remote host.
- Remote management. System managers can monitor the network and applications from remote hosts.
- Remote terminal access. Users can access a remote host as if their terminal is connected directly to that host.
- Remote command execution. Users can issue commands to remote hosts.

Network Considerations

23.3 Understanding Compaq TCP/IP Services for OpenVMS

- Remote printing. Users can send or receive print jobs to or from remote printers.
- Remote file copy. Users can copy files that reside on remote hosts.
- Remote booting. Servers can provide boot information for remote clients.

Internetworking with TCP/IP hides the hardware details of each individual network and allows computers to communicate independently of their physical network connections. TCP/IP provides both a standard transport mechanism and full-duplex, reliable, stream communication services for software applications.

Compaq TCP/IP Services for OpenVMS software provides interoperability and resource sharing between OpenVMS systems, UNIX systems, and other systems that support the TCP/IP protocol suite and Sun Microsystems' Network File System (NFS). TCP/IP systems and other internet hosts share data and resources by using standard TCP/IP protocols over a number of network hardware configurations: Ethernet, Fiber Distributed Data Interface (FDDI), Token Ring, and asynchronous transfer mode (ATM).

Each end system connected to a TCP/IP network is called a **host**. Each host has a unique name and address. The local host is the system you are using, and the remote host is the system with which you are communicating. Hosts are connected by **lines** that carry information between the hosts. The line is the physical path over which data can pass from one host to another. (Examples of lines are telephone lines, fiber-optic cables, and satellites.)

A TCP/IP network is called a packet-switching network. Information is transmitted in small packets of data rather than as a continuous stream from host to host. For example, a file to be transmitted from one host to another is divided into many small packets that are sent across the network one at a time. Each packet contains information about the address of the destination host. At the destination, the packets are reassembled.

The process of directing a data message from a source host to a destination host is called **routing**. For hosts not directly connected to each other, data can be forwarded from the source to the destination through intervening hosts.

23.3.1 Support for OpenVMS Cluster Systems

Compaq TCP/IP Services for OpenVMS supports OpenVMS Cluster systems and the use of cluster aliases. The network sees the cluster as one system with one name, called the internet alias. A remote host can use the cluster alias to address the cluster as one host or use the host name of a cluster member to address a cluster member individually.

23.3.2 TCP/IP Services Management Tools and Utilities

Compaq TCP/IP Services for OpenVMS provides a comprehensive, easy-to-use network management tool that includes over 100 OpenVMS DCL-style commands. These commands allow you to locally configure, monitor, and tune TCP/IP Services components by issuing management commands at the TCPIP> prompt.

You can also use UNIX management commands to manage some components.

For detailed information about managing TCP/IP Services on your system, refer to the *Compaq TCP/IP Services for OpenVMS Management* guide.

Network Considerations

23.4 Preparing to Join a TCP/IP Network

23.4 Preparing to Join a TCP/IP Network

Before you can configure a host to run TCP/IP, you must have a unique IP address and host name. Unlike network hardware addresses, which are hardcoded and fixed, an IP address is assigned by a network administrator. Networks connected to the public Internet must obtain an official, unique network ID from the InterNIC. The IP address has a total of 32 bits (four octets) that identify the network and host. A host name is the name assigned to a computer to facilitate communication.

In addition to providing host and network IDs, the IP address provides information that helps a host determine which packets it should receive and which packets it should ignore. In order for a host to know whether or not another host is on the same or a different subnet, the host compares its own address and the address of the destination computer to a subnet mask. If the network ID of the destination matches the network ID of the source, the packet is delivered to the destination host on the local network. If the network IDs do not match, the packet is forwarded through an IP router to the destination computer.

In addition to an IP address and host name, you must obtain the following information from your network administrator:

- The address for the default gateway, if your systems needs to communicate with TCP/IP hosts not on the local network
- The routing protocol the network uses
- The address of the domain name server (or servers) that will resolve host names into IP addresses
- The network subnet and broadcast masks

You must also decide which end-user services to provide and whether your system is operating as a client, server, or both.

For more information about planning to install and configure Compaq TCP/IP Services for OpenVMS, refer to *Digital TCP/IP Services for OpenVMS Concepts and Planning*.

23.5 Installing and Configuring TCP/IP Services

You can install Compaq TCP/IP Services for OpenVMS software on your system in either of the following ways:

- From the OpenVMS installation procedure menu, as part of the operating system upgrade or installation
- As a layered application using the DCL command `PRODUCT INSTALL`

Once TCP/IP Services software is installed successfully, you invoke the menu-based configuration procedure, as follows, to configure the software according to the unique characteristics of your system and network:

```
$ @SYS$MANAGER:TCPIP$CONFIG
```

For detailed installation and configuration information, refer to the *Compaq TCP/IP Services for OpenVMS Installation and Configuration* manual.

23.6 Starting and Stopping TCP/IP Services

To start Compaq TCP/IP Services for OpenVMS software on your system after you configure the standard software and optional components, or to stop the software for an orderly shutdown, enter:

```
$ @SYS$MANAGER:TCPIP$CONFIG
```

Then enter the appropriate menu option to start or stop TCP/IP Services.

23.7 TCP/IP Services Documentation

Table 23-3 lists the documentation that supports Compaq TCP/IP Services for OpenVMS. For detailed information about how to plan for, install, and use TCP/IP Services, refer to these documents.

Network Considerations

23.7 TCP/IP Services Documentation

Table 23–3 Compaq TCP/IP Services for OpenVMS Documentation

Manual	Contents
<i>Compaq TCP/IP Services for OpenVMS Release Notes</i>	Describes changes to the software including installation, upgrade, and compatibility information. These notes also describe new and existing software problems and restrictions, and software and documentation corrections.
<i>Digital TCP/IP Services for OpenVMS Concepts and Planning</i>	Introduces TCP/IP concepts and components and provides information to help you plan your software configuration.
<i>Compaq TCP/IP Services for OpenVMS Installation and Configuration</i>	Explains how to install and configure the Compaq TCP/IP Services for OpenVMS product on your OpenVMS host.
<i>Digital TCP/IP Services for OpenVMS User's Guide</i>	Describes how to use the applications available with TCP/IP Services such as remote file operations, E-mail, TELNET, TN3270, and network printing. Also explains how to use these services to communicate with systems on private internets or on the worldwide Internet.
<i>Compaq TCP/IP Services for OpenVMS Management</i>	Describes day-to-day management of the Compaq TCP/IP Services for OpenVMS software product.
<i>Compaq TCP/IP Services for OpenVMS Management Command Reference</i>	Describes the Compaq TCP/IP Services for OpenVMS management commands. A companion guide to the <i>Compaq TCP/IP Services for OpenVMS Management</i> manual.
<i>Compaq TCP/IP Services for OpenVMS Tuning and Troubleshooting</i>	Describes how to troubleshoot Compaq TCP/IP Services for OpenVMS and how to tune the performance of the product.
<i>Compaq TCP/IP Services for OpenVMS Guide to IPv6</i>	Describes the Compaq TCP/IP Services for OpenVMS IPv6 features and how to install and configure IPv6 on your system.
<i>Compaq TCP/IP Services for OpenVMS Sockets API and System Services Programming</i>	Describes how to use Compaq TCP/IP Services for OpenVMS to develop network applications using Berkeley Sockets Sockets or OpenVMS system services.
<i>Compaq TCP/IP Services for OpenVMS ONC RPC Programming</i>	Provides an overview of high-level programming with open network computing remote calls (ONC RPC), describes how to use the RPCGEN protocol compiler to create applications, and describes the RPC programming interface.
<i>Compaq TCP/IP Services for OpenVMS SNMP Programming and Reference</i>	Describes the Simple Network Management Protocol (SNMP) and the SNMP application programming interface (eSNMP). It describes the subagents provided with TCP/IP Services, utilities provided for managing subagents, and how to build your own subagents.

23.8 Understanding DECnet-Plus for OpenVMS Networking Software

DECnet-Plus for OpenVMS provides the means for various Compaq operating systems to communicate with each other and with systems provided by other vendors. The DECnet-Plus network supports remote system communication, resource sharing, and distributed processing. Network users can access resources on any system in the network. Each system participating in the network is known as a network **node**.

23.8 Understanding DECnet-Plus for OpenVMS Networking Software

DECnet-Plus is the implementation of the fifth phase of the Digital Network Architecture (DNA). DNA Phase V integrates the OSI protocols with DECnet protocols. In addition, DECnet-Plus includes support for the Internet standard RFC 1006 and the Internet draft RFC 1859, allowing OSI and DECnet applications to run over TCP/IP. Thus, using DECnet-Plus, applications can communicate with peer OSI and DECnet applications on any DECnet Phase IV-based system or OSI-based system, whether from Compaq or from other vendors.

The support of the non-proprietary protocols is the primary difference between DECnet (Phase IV) and DECnet-Plus.

DECnet-Plus provides many features designed to enhance networking capabilities. These features include:

- Global name/directory services that allow large networks to easily store, manage, and access addressing information for (potentially) millions of network objects, such as end systems, users, printers, files, and disks.
- Optional local name/directory services for smaller networks that do not have as critical a need for global directory services.
- Expanded network management capabilities with the Network Control Language (NCL).
- Host-based routing that allows an OpenVMS system to operate as a DECnet-Plus intermediate system in a routing domain.

This feature is especially useful for those configurations where you need to route from a LAN to a WAN and want to use an existing system to do the routing rather than investing in a dedicated router. Host-based routing is not intended for use in network configurations that have high-throughput requirements.

- Increased addressing capacity with the OSI standard address format, making possible unique addressing for virtually an unlimited number of nodes. Existing Phase IV addresses can continue to be used for systems upgraded to DECnet-Plus. The Phase IV address is automatically converted by the configuration procedure to the OSI address format (as a DECnet Phase IV compatible address).
- Address autoconfiguration that enables an adjacent router to configure the node address for the local node.
- Single configuration for OSI components.

X.25 (on VAX systems), wide area device drivers (WANDD), File Transfer, Access, and Management (FTAM), and Virtual Terminal (VT) applications are included with DECnet-Plus software. (On Alpha systems, X.25 support is separate from DECnet-Plus software.)

23.8.1 DECnet-Plus Node Names

Naming conventions for DECnet node names correspond to the two types of DECnet functionality:

- DECnet-Plus full names
Full names are hierarchically structured DECnet node names that can be stored in a DECdns name service. Full names can be a maximum of 255 bytes long.
- DECnet Phase IV node names, called node synonyms in DECnet-Plus

Network Considerations

23.8 Understanding DECnet-Plus for OpenVMS Networking Software

These names are the shorter names used by DECnet Phase IV, restricted to six or fewer characters. Using these names enables DECnet-Plus to be backward compatible with DECnet Phase IV systems in the same network.

Syntax for Full Names

Full names have the following general syntax:

```
namespace:directory ... .directory.node-name
```

where:

<i>namespace</i>	Identifies the global name service
<i>directorydirectory</i>	Defines the hierarchical directory path within the name service
<i>node-name</i>	Is the specific object defining the DECnet node

The following examples show node full names for the Local namespace, DECdns, and DNS/BIND, respectively:

Local namespace	-	LOCAL:.CPlace
DECdns	-	ACME:.warren.CPlace
Domain	-	CPlace.warren.acme.com

The system stores a full name as you enter it, preserving uppercase and lowercase entries. However, when matching an entry with a stored full name, the system is case insensitive; in other words, if the user enters Acme, the system recognizes it as equivalent to ACME.

For more information about full names, refer to the DECnet-Plus documentation.

23.8.2 Support for OpenVMS Cluster Systems

DECnet-Plus software supports OpenVMS Cluster systems and the use of cluster aliases. DECnet-Plus allows for three aliases for each OpenVMS Cluster. DECnet Phase IV nodes cannot be DECnet-Plus alias members. (A separate alias must be configured for use with DECnet Phase IV nodes.)

The CLUSTER_CONFIG.COM command procedure performs an OpenVMS Cluster configuration. It can configure all members of a cluster from any cluster member. It invokes the DECnet-Plus for OpenVMS NET\$CONFIGURE.COM command procedure to perform any required modifications to NCL initialization scripts. Use CLUSTER_CONFIG.COM to configure an OpenVMS Cluster. Use NET\$CONFIGURE.COM directly to configure additional DECnet-Plus satellite nodes once CLUSTER_CONFIG.COM has already been used.

23.8.3 DECnet-Plus Management Tools and Utilities

DECnet-Plus for OpenVMS provides a variety of network tools that let you perform the following tasks:

- Manage local and remote DECnet Phase V components. Two interfaces are available: the Network Control Language (NCL) command-line interface and a Motif-based windows interface (NET\$MGMT).
Compaq supplies the DECNET_MIGRATE tool that converts individual DECnet Phase IV NCP commands to NCL commands, and NCP commands within command procedures to NCL commands. You can use DECNET_MIGRATE when you are new to DECnet-Plus, learning NCL, and want help specifying familiar NCP commands in NCL syntax.
- Manage remote DECnet Phase IV nodes with the NCP Emulator (NCP.EXE). This utility supports a significant range of NCP commands. It is not designed to replace NCL for managing a DECnet-Plus system.

23.8 Understanding DECnet-Plus for OpenVMS Networking Software

- Use DECnet-Plus for OpenVMS initialization scripts (files in the form SYS\$MANAGER:NET\$*.NCL).
- Perform maintenance operations (using MOP) such as downline load, upline dump, remote console connection, and loopback testing support. DECnet-Plus for OpenVMS provides enhanced support and performance for concurrent downline loads. For more information about MOP and how to start this process, refer to the *DECnet-Plus for OpenVMS Network Management* guide.
- Perform enhanced event logging using EVD.
- Use Common Trace Facility (CTF) for troubleshooting.
- Use the DECNET_REGISTER tool to assist in managing node names in your network for the Local and DECdns namespaces.

23.9 Preparing to Join a DECnet-Plus Network

Before configuring your DECnet-Plus node, you must make some decisions regarding addressing, the use of name services, time services, and routers. You must also be aware of license dependencies unique to X.25 software.

For detailed planning information, refer to the *DECnet-Plus Planning Guide*.

23.10 Installing and Configuring DECnet-Plus

You can install DECnet-Plus for OpenVMS software on your system in either of the following ways:

- From the OpenVMS installation procedure menu, as part of the operating system upgrade or installation
- As a layered application using the DCL command PRODUCT INSTALL

Once the DECnet-Plus software is installed successfully, you invoke a menu-based configuration procedure to configure the software according to the unique characteristics of your system and network:

- Use the Fast configuration option if you are upgrading from a DECnet Phase IV node, you plan to use the existing Phase IV configuration, and your node is not part of an OpenVMS Cluster. Enter:

```
$ @SYS$MANAGER:NET$CONFIGURE
```

- Use the Basic configuration option if your node is in a cluster, you are upgrading or reconfiguring DECnet-Plus, and you want to run DECnet over TCP/IP. Enter:

```
$ @SYS$MANAGER:NET$CONFIGURE BASIC
```

- Use the Advanced configuration option if your node's configuration is complex and you need to customize it. Enter:

```
$ @SYS$MANAGER:NET$CONFIGURE ADVANCED
```

For detailed installation and configuration information, refer to the *DECnet-Plus for OpenVMS Installation and Basic Configuration* manual.

Network Considerations

23.11 Using DECnet Over TCP/IP

23.11 Using DECnet Over TCP/IP

The DECnet over TCP/IP feature extends the DECnet Phase V architecture to coexist and communicate with TCP/IP networks. This feature requires a valid DECnet license and a licensed and installed TCP/IP product that supports the PATHWORKS Internet Protocol (PWIP) interface.

You should consider enabling DECnet over TCP/IP for the following purposes:

- Transition from DECnet to TCP/IP. To move to TCP/IP-based networks gradually based on business need. DECnet over TCP/IP functionality is designed to be transparent to users.
- Coexistence. To run DECnet or OSI applications, or both, in a TCP/IP-based network.

With DECnet over TCP/IP, you can:

- Expand your DECnet network to communicate over TCP/IP network backbones. DECnet over TCP/IP allows combinations of naming services (DECdns, large LOCAL file, and DNS/BIND).
- Expand the network by joining two existing DECnet networks without the need to renumber the nodes.
- Use IP-only traffic in part or all of the backbone. With DECnet over TCP/IP, DECnet or OSI applications can run over the lower levels of the TCP/IP protocol stack and the IP network backbone.
- Enable the feature on a node-to-node basis or throughout the entire network.

The DECnet over TCP/IP functionality uses TCP/IP to form a logical link between DECnet nodes. It uses the PATHWORKS IP Driver to interface with TCP. DECnet applications run transparently between DECnet nodes connected by TCP/IP. Users on those DECnet nodes can connect to each other using DECnet node synonyms or IP fullnames. For example:

```
$ SET HOST SYSABC
$ SET HOST SYSABC.boston.acme.com
$ SET HOST 16.12.42.19
```

To enable DECnet over TCP/IP on your system, you must do the following primary tasks:

- Install and configure DECnet-Plus on the system.
When you run the NET\$CONFIGURE procedure, specify the Domain directory service (for TCP/IP addresses) as well as the Local or DNS/BIND service when prompted by the configuration procedure.
- Install and configure TCP/IP Services on the system.
When you run the TCPIP\$CONFIG configuration procedure, enable the PWIP driver to form a bridge between DECnet-Plus and TCP/IP Services software. The PWIP driver is listed as Option 1 on the Optional Components menu.

For detailed information about enabling and using DECnet over TCP/IP, refer to the *DECnet-Plus for OpenVMS Applications Installation and Advanced Configuration* and *DECnet-Plus for OpenVMS Network Management* manuals.

23.12 Moving Your DECnet Phase IV Network to DECnet-Plus

If you are planning to transition your network from DECnet Phase IV to DECnet-Plus, you can move portions of the network to DECnet-Plus or move the entire network. Because DECnet-Plus is backward compatible, you can choose to run your system and the network in the same manner as you have before, using DECnet Phase IV applications, routing, and so forth. You can implement the additional functionality available to you from DECnet-Plus any time you are ready. The changes mostly involve network management. They are almost entirely transparent to users and applications.

A number of automated tools (DECnet transition utilities and the NCP Emulator) are available, in addition to the simplified configuration procedures, to help ease the transition to full DECnet-Plus functionality.

For detailed information about transitioning your network, refer to the *DECnet-Plus Planning Guide* manual.

23.13 Starting and Stopping DECnet-Plus

If you install DECnet-Plus from the OpenVMS installation menu, the software starts automatically. If you need to restart DECnet-Plus for any reason (for example, after shutting down the network by executing `SYS$STARTUP:NET$SHUTDOWN.COM`), enter the following command:

```
$ @SYS$STARTUP:NET$STARTUP
```

To shut down DECnet-Plus software, which disables and deletes various network entities on your system, enter:

```
$ @SYS$MANAGER:NET$SHUTDOWN
```

23.14 DECnet-Plus Documentation

Table 23–4 lists the documentation that supports the DECnet-Plus for OpenVMS software. For complete information about how to plan for, install, configure, and manage DECnet-Plus, refer to these documents.

Table 23–4 DECnet-Plus for OpenVMS Documentation

Manual	Contents
<i>DECnet-Plus for OpenVMS Release Notes</i>	Describes changes to the software; installation, upgrade, and compatibility information; new and existing software problems and restrictions; and software and documentation corrections. You can print this text file from the configuration procedure.
<i>DECnet-Plus for OpenVMS Introduction and User's Guide</i>	Provides an introduction to networking on the system and includes user information.
<i>DECnet-Plus for OpenVMS Installation and Basic Configuration</i>	Explains how to install DECnet-Plus and perform the BASIC configuration option.
<i>DECnet-Plus for OpenVMS Applications Installation and Advanced Configuration</i>	Explains how to install and configure network applications and perform the ADVANCED configuration procedure option.

(continued on next page)

Network Considerations

23.14 DECnet-Plus Documentation

Table 23–4 (Cont.) DECnet-Plus for OpenVMS Documentation

Manual	Contents
<i>DECnet-Plus for OpenVMS Installation Quick Reference Card</i>	Provides a quick reference for upgrading your system to DECnet-Plus during installation.
<i>DECnet-Plus Planning Guide</i>	Provides steps for transitioning from DECnet Phase IV functionality to DECnet Phase V.
<i>DECnet-Plus for OpenVMS Network Management Quick Reference Guide</i>	Provides a quick reference for DECnet-Plus network management.
<i>DECnet-Plus for OpenVMS Network Management</i>	Includes network management concepts and tasks for DECnet-Plus systems.
<i>DECnet-Plus Network Control Language Reference</i>	Provides command descriptions and examples for all NCL commands.
<i>DECnet-Plus Problem Solving</i>	Explains how to isolate and solve DECnet-Plus problems that can occur while the network is running, and how to perform loopback tests.

Managing the Local Area Network (LAN) Software

This chapter describes how the LAN software works and the tasks you must perform to manage the LAN software on your system.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Running the LANACP LAN server process	Section 24.3.1
Invoking and running LANCP	Section 24.4.1
Managing LAN devices	Section 24.5
Managing the LAN device databases	Section 24.6
Managing the LAN node databases	Section 24.7
Migrating from DECnet MOP to LAN MOP	Section 24.8.2
Using CLUSTER_CONFIG_LAN.COM and LAN MOP	Section 24.8.3
Managing the MOP downline load services	Section 24.9
Initiating the MOP console carrier	Section 24.9.8
Requesting MOP trigger boot	Section 24.9.9

This chapter explains the following concepts:

Concept	Section
Local area networks	Section 24.1
LANACP LAN server process	Section 24.3
LANCP utility	Section 24.4
MOP downline load services	Section 24.8

24.1 Understanding Local Area Networks

A local area network (LAN) provides a communications channel designed to connect information processing equipment in a limited area such as a room, a building, or a complex of buildings (for example, a campus). Nodes in a LAN can be linked by the following types of data transmission media:

- **Ethernet**—One of the earliest and the most common LANs. Ethernet can refer to either a general LAN application (for example, Ethernet address) or to the specific CSMA/CD (carrier sense multiple access with collision detection) technology that implements the Intel, Xerox, and Compaq intercompany Ethernet specifications.

Managing the Local Area Network (LAN) Software

24.1 Understanding Local Area Networks

There are three types of Ethernet LANs as follows:

- Ethernet (802.3) with transmission speeds of 10Mbps
- Fast Ethernet (802.3u) with transmission speeds of 100Mbps
- Gigabit Ethernet (802.3z) with transmission speeds of 1000Mbps

All three types of Ethernet employ CSMA/CD protocol, and the same frame format and same frame size.

- FDDI (Fiber Distributed Data Interface) — Implemented as dual-ring, token ring LANs.
- Token Ring—The IEEE 802.5 standard token passing ring.
- ATM (Asynchronous Transfer Mode) — The following standard is supported:
 - LAN emulation over ATM supports the ATM Forum's LAN Emulation V1.0 (LANE) standard.
 - Classical IP over ATM supports the RFC 1577 standard. (DGLTA, DGLPA, DGLPB only)

24.1.1 LAN Characteristics

LAN controllers are devices that, along with additional external hardware, implement the Ethernet, FDDI, Token Ring, LAN emulation over ATM or Classical IP (RFC 1577) specifications. A LAN controller and the local system constitute a node. The LAN controller communicates with the local system through the system bus, and with remote systems implementing the Ethernet, FDDI, Token Ring, or LAN emulation over ATM specifications through the communication medium. (The Ethernet specification is described in *The Ethernet—Data Link Layer and Physical Layer Specification*. The FDDI specifications are available from ANSI. The Token Ring specifications are available from IEEE. The LAN emulation over ATM specifications are available from the ATM Forum.)

Application programs use the LAN driver's QIO interface to perform I/O operations to and from other nodes on the LAN. For detailed information about the QIO interface, refer to the *OpenVMS I/O User's Reference Manual*.

Managing the Local Area Network (LAN) Software

24.1 Understanding Local Area Networks

Table 24–1 provides a brief summary of the differences between the types of LAN media.

Table 24–1 Characteristics of LAN Media

Media	Speed	Maximum Frame Size	Maximum Cable Length†
Ethernet 802.3	10Mbps	1518 bytes	100 meters
Fast Ethernet 802.3u	100Mbps	1518 bytes	100Base-TX—100 meters 100Base-FX—412 meters (one-half duplex) – 2000 meters (full-duplex)
Gigabit Ethernet 802.3z	1000Mbps	1518 or 9018 bytes	1000Base-SX fiber optic – 550 meters 1000Base-LX fiber optic—5 kilometers 1000Base-CX copper shielded—25 meters 1000BaseT copper UTP—100 meters
FDDI	100Mbps	4495 bytes	40 kilometers
Token Ring 802.5	4 or 16 Mbps	4462 bytes	300 meters
LAN emulation over ATM	155Mbps or 622Mbps	1516, 4544, or 9234	2 kilometers, 300 meters

†Larger networks can be constructed with hubs, bridges, and switches.

24.1.1.1 Ethernet LANs

An Ethernet is a cable to which each system or device is connected by a single line. In an office or other area where personal computers and workstations are located, ThinWire Ethernet or unshielded twisted-pair cabling is usually used.

Individual systems can either be connected directly to an Ethernet or gain access to an Ethernet by means of a local area interconnect device, such as a DELNI. A DELNI serves as a concentrator, grouping systems together. Many similar devices, such as hubs, repeaters, and switches also provide the connectivity.

24.1.1.2 FDDI LANs

FDDI uses a dual ring of trees topology. It uses one ring as the primary ring, the other ring as a backup, and the tree configuration for increased flexibility, manageability, and availability.

FDDI networks and Ethernet networks can be combined to form a single extended LAN. This lets applications running on a system connected to FDDI communicate with applications that run on a system connected to Ethernet.

An FDDI concentrator provides for the attachment of FDDI devices such as VAX and Alpha nodes or FDDI-Ethernet bridges to the FDDI LAN.

24.1.1.3 Token Ring LANs (Alpha Only)

Token Ring controllers use either shielded or unshielded twisted pairs of wire to access the ring. Note that it is difficult to connect a Token Ring LAN directly bridged to any other type of LAN. However, routing protocols to other LANs work easily.

Managing the Local Area Network (LAN) Software

24.1 Understanding Local Area Networks

24.1.1.4 ATM LANs (Alpha Only)

LANs over ATM consist of a connection-oriented network based on cell switching. The OpenVMS ATM network uses AAL5 ATM adaptation layer for data transmission.

For LAN emulation over ATM, OpenVMS implements only the LAN emulation client (LEC) and does not implement the LAN emulation server (LES), the Broadcast and Unknown (BUS), or the LAN emulation Configuration Server (LECS). The LES, BUS, and LECS must be provided by some other facility such as the ATM switch. OpenVMS supports eight LAN emulation clients per ATM adapter.

Classical IP over ATM (CLIP) implements a data-link level device that has the same semantics as an Ethernet interface (802.3). This interface is used by a TCP/IP protocol to transmit 802.3 (IEEE Ethernet) frames over an ATM network. The model that OpenVMS follows for exchanging IP datagrams over ATM is based on RFC1577 (Classical IP over ATM).

24.1.2 LAN Addresses

Nodes on the LAN are identified by unique addresses. A message can be sent to one, several, or all nodes on the LAN simultaneously, depending on the address used.

Upon application, IEEE assigns a block of addresses to a producer of LAN nodes. Thus, every manufacturer has a unique set of addresses to use. Normally, one address out of the assigned block of physical addresses is permanently associated with each controller (usually in read-only memory). This address is known as the hardware address of the controller. Each controller has a unique hardware address.

A LAN address is 48 bits in length. LAN addresses are represented as six pairs of hexadecimal digits (six bytes) separated by hyphens (for example, AA-01-23-45-67-FF). The bytes are displayed from left to right in the order in which they are transmitted; bits within each byte are transmitted from right to left. In this example, byte AA is transmitted first; byte FF is transmitted last.

A LAN address can be a physical address of a single node or a multicast address, depending on the value of the low-order bit of the first byte of the address (this bit is transmitted first). The two types of node addresses are:

- Physical address—The unique address of a single node on a LAN. The least significant bit of the first byte of a physical address is 0. (For example, in physical address AA-00-03-00-FC-00, byte AA in binary is 1010 1010, and the value of the low-order bit is 0.)
- Multicast address—A multidestination address of one or more nodes on a given LAN. The least significant bit of the first byte of a multicast address is 1. (For example, in the multicast address 0B-22-22-22-22-22, byte 0B in binary is 0000 1011, and the value of the low-order bit is 1.)

Token Ring devices do not support IEEE 802 standard multicast addresses. They do support functional addresses. A functional address is a locally administered group address that has 31 possible values. Each functional address sets one bit in the third through sixth bytes of the address, and bytes 1 and 2 are 03-00 (C0:00 in bit reversed format). To convert a multicast address to a functional address, use the SET DEVICE/MAP command.

24.2 Managing Local Area Networks

The local area network (LAN) software includes two system management tools that work in conjunction with the OpenVMS LAN driver system software:

- Local Area Network Control Program (LANCP)
- LAN Auxiliary Control Program (LANACP) LAN server process

The LAN system management tools:

- Allow you to set LAN parameters to customize your LAN environment.
- Display LAN settings and counters.
- Provide Maintenance Operations Protocol (MOP) downline load support for devices such as terminal servers, x-terminals, and LAN-based printers, and for booting satellites in an OpenVMS Cluster environment. This MOP support provides an alternative to the traditional method of using either DECnet for OpenVMS or DECnet/OSI software.

Table 24–2 describes the LAN management software and the functionality supported on systems running OpenVMS Alpha and OpenVMS VAX.

Table 24–2 LAN System Management Enhancements

Utility	Description	OpenVMS Support
LAN Auxiliary Control Program (LANACP)	Runs as a server process whose primary function is to provide MOP downline load service. Other services include maintenance of a LAN volatile device database and a LAN volatile node database.	The LANACP utility provides identical functionality on VAX and Alpha systems running OpenVMS Version 7.0 and later.

(continued on next page)

Managing the Local Area Network (LAN) Software

24.2 Managing Local Area Networks

Table 24–2 (Cont.) LAN System Management Enhancements

Utility	Description	OpenVMS Support			
LAN Control Program (LANCP)	<p>Allows you to control LAN software parameters and obtain information from the LAN software. You can use the LANCP utility to:</p> <ul style="list-style-type: none"> Obtain LAN device counters, revision, and configuration information Change the operational parameters of LAN devices on the system Maintain the permanent and volatile LAN device and node databases Control the LANACP LAN server process (including MOP downline load server related functions) Initiate MOP console carrier connections Send MOP trigger boot requests to other nodes 	OpenVMS Alpha Version 6.1 contained the initial implementation of LANCP, which did not include MOP-related functions.			
		OpenVMS Version 6.2 (VAX and Alpha) added MOP-related functions and extended some of this capability to VAX systems. The following table shows how the LAN utility functions are currently supported on VAX and Alpha systems:			
		Function	OpenVMS Alpha Version 7.3-1	OpenVMS VAX Version 7.3	
		Change operational parameters of LAN devices?	Yes	No	
		Display LAN device information?	Yes	Limited	
		Support MOP functions?	Yes	Yes	

24.3 Understanding the LANACP LAN Server Process

You can run the LANACP LAN server process to provide the following services:

- Maintenance of the LAN volatile node database
- Maintenance of the LAN volatile device database
- MOP downline load

The LANCP utility allows you to issue instructions to the LANACP process.

Three principal files are connected with LANACP:

- SYS\$SYSTEM:LANACP.EXE**
This file is the LANACP utility program.
- SYS\$STARTUP:LAN\$STARTUP.COM**
This file starts the LANACP server process.
- SYS\$STARTUP:VMS\$DEVICE_STARTUP.COM**
This file contains an entry that is used to start LANACP automatically at system startup.

In addition, four system logical names, described in Table 24–3, are associated with the LANACP LAN server process.

Managing the Local Area Network (LAN) Software

24.3 Understanding the LANACP LAN Server Process

Table 24–3 LANACP System Logical Names

Component	Description
LAN\$DLL	Defines the location of downline load files, where the location of the file is not provided in the load request or explicitly defined in the LAN volatile node database. By default, this is defined as SYS\$SYSROOT:[MOM\$SYSTEM].
LAN\$NODE_DATABASE	Defines the location of the LAN permanent node database. By default, this is defined as SYS\$COMMON:[SYSEXE]LAN\$NODE_DATABASE.DAT.
LAN\$DEVICE_DATABASE	Defines the location of the LAN permanent device database. By default, this is defined as SYS\$SPECIFIC:[SYSEXE]LAN\$DEVICE_DATABASE.DAT.
LAN\$ACP	Defines the location of the LANACP LAN server process log file, containing entries describing changes to the LAN permanent device and node databases, and load request and load status information. By default, this is defined as SYS\$MANAGER:LAN\$ACP.LOG.

24.3.1 Running the LANACP LAN Server Process

To start the LANACP LAN server process, type @SYS\$STARTUP:LAN\$STARTUP at the DCL prompt.

24.3.2 Stopping the LANACP LAN Server Process

To stop the LANACP LAN server process, enter the SET ACP/STOP command at the LANCP utility prompt.

24.4 Understanding the LANCP Utility

The LANCP utility allows you to set and show LAN parameters. Section 24.4.1 describes how to invoke the LANCP utility. Table 24–4 describes LAN functions and provides section references to the LANCP commands that help you perform these functions.

Table 24–4 Functions of the LANCP Utility

Task	Section
Managing LAN devices	Section 24.5
Managing LAN device databases	Section 24.6
Managing LAN node databases	Section 24.7
Managing the MOP downline load service	Section 24.9
Initiating a MOP console carrier connection	Section 24.9.8
Sending MOP trigger boot requests	Section 24.9.9

24.4.1 Invoking and Running LANCP

Table 24–5 describes the ways you can invoke and run the LANCP utility (SYS\$SYSTEM:LANCP.EXE).

Managing the Local Area Network (LAN) Software

24.4 Understanding the LANCP Utility

Table 24–5 Invoking the LANCP Utility

Command	Example
Use the RUN command	At the DCL command prompt, enter: \$ RUN SYSS\$SYSTEM:LANCP The LANCP utility displays the LANCP prompt at which you can enter LANCP commands.
Define LANCP as a foreign command	Either at the DCL prompt or in a startup or login command file, enter: \$ LANCP ::= SYSS\$SYSTEM:LANCP Then, you can enter the command LANCP at the DCL prompt to invoke the utility and enter LANCP commands. When you enter the LANCP command: <ul style="list-style-type: none"> • Without specifying any command qualifiers, the LANCP utility displays the LANCP prompt at which you can enter commands. • With command qualifiers, the LANCP utility terminates after it executes the command and returns you to the DCL prompt.
Use the MCR command	At the DCL command prompt, enter: \$ MCR LANCP When you enter the MCR LANCP command: <ul style="list-style-type: none"> • Without specifying any command qualifiers, the LANCP utility displays the LANCP prompt at which you can enter commands. • With command qualifiers, the LANCP utility terminates after it executes the command and returns you to the DCL prompt.

At the LANCP> prompt, you can enter LANCP commands.

For information about the LANCP utility, enter the HELP command at the LANCP> prompt.

To exit from the LANCP utility, enter the EXIT command or press Ctrl/Z at the LANCP> prompt.

24.4.2 LANCP Commands

Table 24–6 summarizes the LANCP commands.

Table 24–6 LANCP Commands

Command	Function
@ (Execute Procedure)	Executes a command procedure.
CLEAR DEVICE	Deletes a device from the LAN volatile device database.
CLEAR DLL	Clears MOP downline load counters for all nodes and devices.
CLEAR MOPDLL	Same as the CLEAR DLL command.
CLEAR NODE	Deletes a node from the LAN volatile node database.

(continued on next page)

Managing the Local Area Network (LAN) Software

24.4 Understanding the LANCP Utility

Table 24–6 (Cont.) LANCP Commands

Command	Function
CONNECT NODE	Connects to a LAN device, such as a terminal server, that implements a management interface using the MOP console carrier protocol.
CONVERT DEVICE_DATABASE	Converts a device database to the current format required by LANCP.
CONVERT NODE_DATABASE	Converts a node database to the current format required by LANCP.
DEFINE DEVICE	Enters a device into the LAN permanent device database or modifies an existing entry.
DEFINE NODE	Enters a node into the LAN permanent node database or modifies an existing entry.
EXIT	Stops execution of LANCP and returns control to the DCL command level.
HELP	Provides online help information about the LANCP utility.
LIST DEVICE	Displays information in the LAN permanent device database.
LIST NODE	Displays information in the LAN permanent node database.
PURGE DEVICE	Deletes a device from the LAN permanent device database.
PURGE NODE	Deletes a node from the LAN permanent node database.
SET ACP	Modifies the operation of the LANACP LAN server process.
SET DEVICE	Modifies device characteristics in the LAN volatile device database and in the device itself.
SET NODE	Enters a node into the LAN volatile node database or modifies an existing entry.
SHOW CONFIGURATION	Displays a list of LAN devices on the system.
SHOW DEVICE	Displays information in the LAN volatile device database.
SHOW DLL	Displays the current state of MOP downline load services.
SHOW LOG	Displays recent downline load activity.
SHOW MOPDLL	Same as the SHOW DLL command.
SHOW NODE	Displays information in the LAN volatile node database.
SPAWN	Creates a subprocess of the current process.
TRIGGER NODE	Issues a request to reboot to a remote node.

For detailed information about LANCP commands and qualifiers, refer to the *OpenVMS System Management Utilities Reference Manual: A–L*.

Managing the Local Area Network (LAN) Software

24.4 Understanding the LANCP Utility

24.4.3 LANCP Miscellaneous Functions

Use the SPAWN command to create a subprocess of the current process. The SPAWN command copies the context of the subprocess from the current process. This allows you to exit temporarily from LANCP without having to restart LANCP when you resume.

The syntax for the SPAWN command is as follows:

```
SPAWN [command-string]
```

You can set up the LANCP utility to execute commands from a command file from within LANCP. The LANCP utility recognizes the command file as the file name preceded by the at sign (@). The default file name extension is .COM.

24.5 Managing LAN Devices

LAN device management consists of displaying device characteristics and setting device parameters. You can use the LANCP utility to set parameters for the types of LAN devices shown in Table 24–7.

Table 24–7 LAN Devices

LAN	Device Examples	Description
Ethernet	DE425, DE434, DE435, DE436, DE500, DECchip 21040, DEMNA	Allow the selection of media type (type of cable connected) and the speed of connection (Ethernet or FastEthernet). Allow full-duplex operation (point-to-point operation between a similar device or between the device and a switch).
FDDI	DEFTA, DEFPA, DEFAA, DEFEA, DEMFA	Allow full-duplex operation.
Token Ring‡	DETRA, DW300, DW110, TC4048	Allow the setting of Token Ring parameters and the definition of source routing and functional address mapping.
All	Any	Allow the setting of generic parameters such as the number of receive buffers.
ATM‡	DGLTA, DGLPA, DGLPB, DAPBA, DAPCA	Allow the setting of Emulated LAN (ELAN) parameters.

‡Alpha only

24.5.1 Displaying System Devices

To display the LAN devices on the system, enter the SHOW CONFIGURATION command using the following syntax:

```
SHOW CONFIGURATION
```

Example

```
LANCP> SHOW CONFIGURATION
```

```
LAN Configuration:
```

Device	Medium	Default LAN Address	Version
EWA0	CSMA/CD	08-00-2B-E4-00-BF	02000023
EWB0	CSMA/CD	08-00-2B-92-A4-0D	02000023
IRA0	Token Ring	00-00-93-58-5D-32	20000223

This command displays the output from a SHOW CONFIGURATION command that was entered on a node that has three LAN devices: two DE435s, and a DETRA:

The version is the device-specific representation of the actual version. In this example, for two devices on the PCI bus, the actual version is in the low byte (2.3 for the DE435 adapters). A device that does not have a readable version is shown as version zero.

Consult your device-specific documentation to correlate the version returned with a particular hardware or firmware implementation of the device.

24.5.2 Displaying Device Characteristics

To display information about a LAN device (in the volatile device database), enter the LANCP command `SHOW DEVICE` using the following syntax:

```
SHOW DEVICE device-name [qualifiers]
```

Table 24–8 provides a brief description of the `SHOW DEVICE` command qualifiers.

Note

If you do not specify a qualifier, the utility displays the matching devices without additional information.

Table 24–8 SHOW DEVICE Command Qualifiers

Qualifier	Description
<code>/ALL</code>	Shows all devices which match device name.
<code>/CHARACTERISTICS</code>	Same as the <code>/PARAMETERS</code> qualifier.
<code>/DLL</code>	It shows downline load characteristics.
<code>/COUNTERS</code>	Displays device counters.
<code>/INTERNAL_COUNTERS</code>	Displays internal counters. By default, it does not display zero counters. To see all counters, including zero, use the additional qualifer <code>/ZERO</code> . To see debug counters, use the additional qualifier <code>/DEBUG</code> .
<code>/MAP</code>	Displays the current configuration of the functional address mapping table.
<code>/MOPDLL</code>	Same as the <code>/DLL</code> qualifier.
<code>/OUTPUT‡</code>	Directs output to the specified filename.
<code>/PARAMETERS</code>	Displays status and related information about the device.
<code>/REVISION</code>	Displays the current firmware revision of the adapter, if available or applicable.
<code>/SR_ENTRY</code>	Displays the contents of the current source routing cache table.
<code>/TRACE</code>	Displays LAN driver trace data.

‡Alpha only

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

Examples

1.

```
LANCP> SHOW DEVICE/COUNTERS EXA0
Device Counters EXA0:
      Value Counter
      ---- -
      259225 Seconds since last zeroed
      5890496 Data blocks received
      4801439 Multicast blocks received
      131074 Receive failure
      764348985 Bytes received
      543019961 Multicast bytes received
      3 Data overrun
      1533610 Data blocks sent
      115568 Multicast packets transmitted
      122578 Blocks sent, multiple collisions
      86000 Blocks sent, single collision
      189039 Blocks sent, initially deferred
      198120720 Bytes sent
      13232578 Multicast bytes transmitted
      7274529 Send failure
      0 Collision detect check failure
      0 Unrecognized frame destination
      0 System buffer unavailable
      0 User buffer unavailable
```

This SHOW DEVICE command displays counters for Ethernet device EXA0.

2.

```
LANCP> SHOW DEVICE/MAP ICA0
Multicast to Functional Address Mapping ICA0:
Multicast address      Functional Address      Bit-Reversed
-----
09-00-2B-00-00-04      03-00-00-00-02-00      C0:00:00:00:40:00
09-00-2B-00-00-05      03-00-00-00-01-00      C0:00:00:00:80:00
CF-00-00-00-00-00      03-00-00-08-00-00      C0:00:00:10:00:00
AB-00-00-01-00-00      03-00-02-00-00-00      C0:00:40:00:00:00
AB-00-00-02-00-00      03-00-04-00-00-00      C0:00:20:00:00:00
AB-00-00-03-00-00      03-00-08-00-00-00      C0:00:10:00:00:00
09-00-2B-02-00-00      03-00-08-00-00-00      C0:00:10:00:00:00
09-00-2B-02-01-0A      03-00-08-00-00-00      C0:00:10:00:00:00
AB-00-00-04-00-00      03-00-10-00-00-00      C0:00:08:00:00:00
09-00-2B-02-01-0B      03-00-10-00-00-00      C0:00:08:00:00:00
09-00-2B-00-00-07      03-00-20-00-00-00      C0:00:04:00:00:00
09-00-2B-00-00-0F      03-00-40-00-00-00      C0:00:02:00:00:00
09-00-2B-02-01-04      03-00-80-00-00-00      C0:00:01:00:00:00
09-00-2B-02-01-07      03-00-00-02-00-00      C0:00:00:40:00:00
09-00-2B-04-00-00      03-00-00-04-00-00      C0:00:00:20:00:00
09-00-2B-02-01-00      03-00-00-00-08-00      C0:00:00:00:10:00
09-00-2B-02-01-01      03-00-00-00-10-00      C0:00:00:00:08:00
09-00-2B-02-01-02      03-00-00-00-20-00      C0:00:00:00:04:00
03-00-00-00-00-01      03-00-00-00-00-01      C0:00:00:00:00:80
03-00-02-00-00-00      03-00-02-00-00-00      C0:00:40:00:00:00
```

This SHOW DEVICE command displays mapping information for Token Ring device ICA0.

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

3.

```
LANCP> SHOW DEVICE/PARAM IRA0
Device Parameters IRA0:
      Value  Parameter
      ----  -
      Normal Controller mode
      External Internal loopback mode
00-00-93-58-5D-32 Hardware LAN address
      Token Ring Communication medium
      Enabled Functional address mode
      No Full duplex enable
      No Full duplex operational
      16 Line speed (megabits/second)
      16 Mbps Ring speed
      STP Line media
      Enabled Early token release
      Disabled Monitor contender
      200 SR cache entries
      2 SR discovery timer
      60 SR Aging Timer
      Enabled Source routing
      3 Authorized access priority
AA-00-04-00-92-FF Upstream neighbor
      0 Ring number
```

This SHOW DEVICE command displays status and parameter information for Token Ring device IRA0.

4.

```
LANCP> SHOW DEVICE/REVISION FXA0
Device revision FXA0: 05140823
```

This command displays revision information for FDDI device FXA0.

5.

```
LANCP> SHOW DEVICE/SR_ENTRY ICA0
Source Routing Cache Table ICA0:
      LAN address      State      XmtTmo      RcvTmo      StaleTmo      DiscvTmo
      -----
AA-00-04-00-92-FF  LOCAL      00000028    00000028    00000245    00000000
```

This SHOW DEVICE command displays source routing entry information for Token Ring device ICA0.

24.5.3 Setting Device Characteristics

All LAN devices are characterized by a collection of parameters. The parameters define the operational characteristics of a LAN device on the medium to which the device is connected.

To set LAN device parameters directly, enter the SET DEVICE command at the LANCP> prompt. The LANCP utility issues this command to the LANACP server process, which then issues the appropriate QIOs to set the driver characteristics.

The syntax for the SET DEVICE command is:

```
SET DEVICE device-name [/qualifiers]
```

Table 24–9 provides a brief description of the SET DEVICE command qualifiers that apply directly to LAN devices.

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

Table 24–9 DEFINE DEVICE and SET DEVICE Command Qualifiers

Qualifier	Description
<code>/AGING_TIMER=<i>value</i></code>	Sets the amount of time in seconds to age source routing cache entries before marking them stale.
<code>/ALL</code>	Sets data for all LAN devices.
<code>/ATMADDRESS=LES‡</code>	<p>Sets the LAN emulation server (LES) address for asynchronous transfer mode (ATM). Usually the address is not user specified, and this qualifier is used only if you want a specific address. By default the address is determined by software from the configuration server for the LES.</p> <p>The <code>/ATMADDRESS=LES</code> qualifier's syntax is as follows:</p> <pre>SET DEVICE/ATMADDRESS = ([NO]LES=the ATM server)</pre>
<code>/ATMADDRESS=ARP‡</code>	<p>Sets the address resolution protocol (ARP) server address for Classical IP over ATM. This qualifier is required before a LIS is enabled if the local host is not the ARP server.</p> <p>The <code>/ATMADDRESS=ARP</code> qualifier's syntax is as follows:</p> <pre>SET DEVICE/ATMADDRESS = (ARP=atm_arp_server)</pre>
<code>/AUTONEGOTIATE</code> (default) <code>/NOAUTONEGOTIATE</code>	Enables or disables the use of auto-negotiation to determine the link settings. You may need to disable link auto-negotiation when connected to a switch or device that does not support auto-negotiation.
<code>/CACHE_ENTRIES=<i>value</i></code>	Sets the number of entries to reserve for caching source routing address entries.

‡Alpha only

(continued on next page)

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

Table 24–9 (Cont.) DEFINE DEVICE and SET DEVICE Command Qualifiers

Qualifier	Description								
/CLIP‡	<p>Sets up the Classical Internet Protocol (CLIP) over ATM (RFC1577). The CLIP qualifier implements a data-link level device as a client and/or a server in a logical IP subnet (LIS). This allows the IP protocol to transmit Ethernet frames over the ATM network. The /CLIP = ENABLE command causes the the system to join the LIS. The /CLIP = DISABLE command causes the client to leave the logical IP subnet.</p> <p>Note that a LIS requires a server, and there must be only one server for each subnet. Communication between subnets can only be performed by a router. There can only be one client for each ATM adapter.</p> <p>The /CLIP qualifier's syntax with standard Internet dotted notation is as follows:</p> <pre>SET DEVICE/CLIP =(ip_subnet=a.b.c.d, ip_address=a.b.c.d, parent=device, name="ip subnet name", enable, disable type = client server)</pre> <p>The meanings of the syntax for /CLIP are as follows:</p> <ul style="list-style-type: none"> • <i>ip_address</i> Specifies the IP address of the CLIP client. • <i>subnet_mask</i> Specifies the subnet mask of the CLIP client. • <i>parent</i> Parent device name. • <i>name</i> Specifies a name for the LIS to aid in operations and diagnostics. • <i>type=client</i> Starts up a classical IP client only. This is the default. • <i>type=server</i> Starts up a classical IP server. Only one server for each LIS is allowed, and the server needs to be started first. • <i>type=(server,client)</i> Starts up a classical IP server and client. <p>Keywords and their meanings for /CLIP are as follows:</p> <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left;">Keyword</th> <th style="text-align: left;">Meaning</th> </tr> </thead> <tbody> <tr> <td>Create</td> <td>Loads the classical IP driver but does not start it.</td> </tr> <tr> <td>Enable</td> <td>Causes the node to join the logical IP subnet.</td> </tr> <tr> <td>Disable</td> <td>Causes the node to leave the logical IP subnet.</td> </tr> </tbody> </table>	Keyword	Meaning	Create	Loads the classical IP driver but does not start it.	Enable	Causes the node to join the logical IP subnet.	Disable	Causes the node to leave the logical IP subnet.
Keyword	Meaning								
Create	Loads the classical IP driver but does not start it.								
Enable	Causes the node to join the logical IP subnet.								
Disable	Causes the node to leave the logical IP subnet.								

‡Alpha only

(continued on next page)

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

Table 24–9 (Cont.) DEFINE DEVICE and SET DEVICE Command Qualifiers

Qualifier	Description
/CONTENDER /NOCONTENDER	Specifies that the Token Ring device is to participate in the Monitor Contention process when it joins the ring. The /NOCONTENDER qualifier directs the device not to challenge the current ring server.
/DEVICE_SPECIFIC= FUNCTION="xxxx", VALUE=n	<p>Allows some device-specific parameters to be adjusted. The following list shows the commands and their meanings.</p> <p>FUNCTION="CCOU"—Clears all device and driver counters. If the value is supplied, it is ignored.</p> <p>FUNCTION="DXMT",VALUE=n—Changes the transmit delay value which is the number of microseconds after completion of a transmit request that an interrupt is generated. The current setting is displayed in the internal counters. This function is applicable to Gigabit NICs.</p> <p>FUNCTION="DRCV",VALUE=n—Changes the receive delay value which is the number of microseconds after completion of a receive that an interrupt is generated. The current setting is displayed in the internal counters. This function is applicable to Gigabit NICs.</p> <p>FUNCTION="CXMT",VALUE=n—Changes the transmit coalesce value which is the number of transmit buffer descriptors that are processed before an interrupt is generated. An interrupt may be generated earlier if transmit delay threshold is reached or when an interrupt on behalf of receive or a link state change is generated. The current setting is displayed in the internal counters. This function is applicable to Gigabit NICs.</p> <p>FUNCTION="CRCV",VALUE=n—Changes the receive coalesce value which is the number of receive buffer descriptors that are filled in before an interrupt is generated. An interrupt may be generated earlier if receive delay threshold is reached or when an interrupt on behalf of transmit or a link state change is generated. The current setting is displayed in the internal counters. This function is applicable to Gigabit NICs.</p>
/DISCOVERY_ TIMER=value	Sets the number of seconds to wait for a reply from a remote node when performing the source routing route discovery process.

(continued on next page)

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

Table 24–9 (Cont.) DEFINE DEVICE and SET DEVICE Command Qualifiers

Qualifier	Description
<i>/DLL=(enable-option, exclusive-option, size-option, knownclientsonly-option)</i>	<p>Provides the MOP downline load service settings for the device.</p> <p>In this qualifier, you can specify:</p> <ul style="list-style-type: none">• <i>enable-option</i> Indicates that MOP downline load service should be enabled or disabled for the device.• <i>exclusive-option</i> Indicates that no other provider of MOP downline load service is allowed on the specified LAN device at the same time as LANACP.• <i>knownclientsonly-option</i> Indicates that MOP downline load requests should be serviced only for clients defined in the LAN volatile node database.• <i>size-option</i> Specifies the size in bytes of the file data portion of each downline load message.
<i>/EARLY</i>	Enables Early Token Release on the device.

(continued on next page)

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

Table 24–9 (Cont.) DEFINE DEVICE and SET DEVICE Command Qualifiers

Qualifier	Description
/ELAN‡	<p>Sets LAN emulation. The /ELAN qualifier has two values: enable and disable. With /ELAN=ENABLE with the keyword STARTUP, the LAN emulation is loaded when LANACP starts. With /ELAN=DISABLE, the same parameters as ENABLE can be used.</p> <p>The /ELAN qualifier's syntax is as follows:</p> <pre>SET DEVICE/ELAN =(parent=parent device, name="ELAN NAME to join", size=1516 type=CSMACD Enable, Disable, description = "description string,")</pre>

The meaning of the syntax for /ELAN are as follows:

- *parent*
The ATM adapter device name. An example of the parent device for DGLTA is: HC*n*0, where *n* is the controller number. An example of the parent device for DAPCA is: HW*n*0, where *n* is the controller number.
- *name*
Optionally specified if you want to join a specific ELAN. The default is null.
- *size*
Maximum frame size of the LAN you want to join. Valid sizes are 1516, 4544, or 9234. The default is 1516.
- *type*
Support currently only for CSMACD, which is the default.
- *description*
A method of describing the ELAN for display purposes only.

Keywords and their meanings for /ELAN are as follows:

Keyword	Meaning
Create	Loads the emulation driver SYS\$ELDRIVER.EXE, but does not start it.
Enable	Begins a join on a specified emulated LAN. It also loads the driver, if not already loaded.
Disable	Causes a client to leave the emulated LAN.

‡Alpha only

(continued on next page)

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

Table 24–9 (Cont.) DEFINE DEVICE and SET DEVICE Command Qualifiers

Qualifier	Description
<code>/FULL_DUPLEX</code> <code>/NOFULL_DUPLEX</code>	Enables full-duplex operation of a LAN device. Before full-duplex operation results from the use of this qualifier, additional device or network hardware setup may be required. Some devices may be enabled for full-duplex operation by default. Some devices may not allow the setting to be changed. The <code>/NOFULL_DUPLEX</code> qualifier disables full-duplex operations.
<code>/JUMBO</code> <code>/NOJUMBO</code> (default)	Enables the use of jumbo frames on a LAN device. Only the Gigabit Ethernet NICs support jumbo frames.
<code>/MAP=(MULTICAST_ ADDRESS=address, FUNCTIONAL_ ADDRESS=address)</code>	Defines a functional address mapping entry.
<code>/MAX_BUFFERS=value</code>	Sets the maximum number of receive buffers to be allocated and used by the LAN driver for the LAN device.
<code>/MEDIA=value</code>	<ul style="list-style-type: none"> • For Token Ring devices: Selects the type of cable that is being used to connect the adapter to the Token Ring Media Access Unit (MAU) for devices that do not automatically detect this. • For Ethernet devices: Selects the cable connection.
<code>/MIN_BUFFERS=value</code>	Sets the minimum number of receive buffers to be allocated and used by the LAN driver for the LAN device.
<code>/PERMANENT_ DATABASE (SET command only)</code>	Updates the device entries in the LAN volatile device database with any data currently set in the permanent database.
<code>/PVC=(vci[,...])‡</code>	Defines the permanent virtual circuit (PVC). This is an optional qualifier. A list of PVCs is defined for use by CLIP clients. This command should be used before enabling the CLIP client. PVC has to be setup manually in the ATM switch. The vci is the VCI (Virtual Circuit ID) of the PVC.
<code>/NOPVC=(vci[,...])‡</code>	Does not set the permanent virtual circuit (PVC).
<code>/RING_PURGER‡</code>	Enables the ring purging process of the FDDI device.
<code>/SOURCE_ROUTING‡</code>	Enables source routing on the Token Ring device.
<code>/SPEED=value</code>	Sets the speed of the LAN, if multiple speeds are supported.
<code>/SR_ENTRY=(LAN_ ADDRESS=address, RI=routing-information)</code>	Statically defines a specific source-routed route for a specific node.
<code>/TOKEN_ROTATION‡</code>	Sets the requested token rotation time for the FDDI ring.

‡Alpha only

(continued on next page)

Managing the Local Area Network (LAN) Software

24.5 Managing LAN Devices

Table 24–9 (Cont.) DEFINE DEVICE and SET DEVICE Command Qualifiers

Qualifier	Description
/TOKEN_TIMEOUT‡	Sets the restricted token timeout time for the FDDI ring.
/TRANSMIT_TIMEOUT‡	Sets the valid transmission time for the FDDI device.
/UPDATE	Adds LAN devices that are not currently in one of the LAN device databases to that database. The DEFINE DEVICE command applies to the permanent database; the SET DEVICE command applies to the volatile database.
/VOLATILE_DATABASE (DEFINE command only)	Updates the device entries in the LAN permanent device database with any data currently set in the volatile database.

‡Alpha only

Examples

1.

```
LANCP> SET DEVICE/CONTENDER/MEDIA=UTP/NOEARLY/SOURCE ICA0
```

This command enables monitor contention, UTP cable media, and source routing, and disables early token release for Token Ring device ICA0.

2.

```
LANCP> SET DEVICE/MEDIA=TWIST EWB0
```

This command sets the media type to twisted pair for the second Tulip Ethernet device.

3.

```
LANCP> SET DEVICE/ALL/MIN_BUFFERS=12
```

This command sets the number of receive buffers for all LAN devices to be no less than 12.

4.

```
LANCP> DEFINE DEVICE EXA0/MOPDLL=(ENABLE,EXCLUSIVE)
```

This command defines LAN device EXA0 to enable LANACP MOP downline load service in exclusive mode. The settings of the KNOWNCLIENTSONLY and SIZE characteristics are not changed. If the device entry does not currently exist in the LAN permanent device database, these settings will be set to the defaults.

5.

```
LANCP> DEFINE DEVICE/ALL/MOPDLL=NOEXCLUSIVE
```

This command sets all LAN devices defined in the LAN permanent device database to nonexclusive mode for LANACP MOP downline load service.

6.

```
LANCP> SET DEVICE EXA0/MOPDLL=(ENABLE,NOEXCLUSIVE)  
LANCP> SET DEVICE FXA0/MOPDLL=(ENABLE,EXCL,KNOWN)
```

These commands enable LANACP MOP downline load service for:

- LAN device EXA0 in nonexclusive mode
- LAN device FXA0 in exclusive mode for only known clients

24.6 Managing the LAN Device Databases

The LAN volatile and permanent device databases contain a single entry for each LAN device that exists on the system. Each entry in the LAN volatile device database contains device information and MOP downline load counters information. Each entry in the LAN permanent device database contains device information that is used to populate the volatile database when the LANACP LAN server process is started.

Typically, each database contains the same devices. However, the permanent database may contain entries for devices that have not yet been configured or installed in the system. The LANACP LAN server process maintains the volatile device database. The LANCP utility maintains the permanent device database. You can manipulate either database using the LANCP utility commands depending on your user privileges, as follows:

- Privileged users can add or delete device entries from each database, enable or disable MOP downline load service, and clear MOP downline load counters information for LAN devices.
- Unprivileged users can view the MOP downline load status and counters information.

The following sections describe how to enter and remove devices from the LAN permanent and volatile device databases, and how to enable and disable MOP downline load services.

24.6.1 Displaying Devices in the LAN Device Databases

To display information in the LAN permanent device database, enter the LIST DEVICE command using the following syntax:

```
LIST DEVICE device-name [/qualifiers]
```

To display information in the LAN volatile device database, enter the SHOW DEVICE command using the following syntax:

```
SHOW DEVICE device-name [/qualifiers]
```

Table 24–10 provides a brief description of the LIST DEVICE and SHOW DEVICE qualifiers.

Table 24–10 LIST DEVICE and SHOW DEVICE Command Qualifiers

Qualifier	Description
/ALL	Lists, or shows all devices which match device names.
/CHARACTERISTICS	Same as the /PARAMETER qualifier.
/COUNTERS†	Displays device counters.
/DLL	Lists or shows downline load characteristics.
/MAP	Displays the current configuration of the functional address mapping table.
/MOPDLL	Same as DLL.
/OUTPUT= <i>file-name</i>	Creates the specified file and directs output to it.

†SHOW DEVICE only

(continued on next page)

Managing the Local Area Network (LAN) Software

24.6 Managing the LAN Device Databases

Table 24–10 (Cont.) LIST DEVICE and SHOW DEVICE Command Qualifiers

Qualifier	Description
/PARAMETERS	Displays status and related information about the device.
/REVISION†	Displays the current firmware revision of the adapter, if available or applicable.
/SR_ENTRY	Displays the contents of the current source routing cache table.
†SHOW DEVICE only	

Note

If you do not specify a qualifier, the utility displays the matching devices without additional information.

24.6.2 Entering Devices into the LAN Device Databases

To enter a device into the LAN permanent device database or to modify an existing entry, enter the DEFINE DEVICE command using the following syntax:

```
DEFINE DEVICE device-name [/qualifiers]
```

To enter a device into the LAN volatile device database or to modify an existing entry, enter the SET DEVICE command using the following syntax:

```
SET DEVICE device-name [/qualifiers]
```

24.6.3 Deleting Devices from the LAN Device Databases

To delete a device from the LAN permanent device database, enter the PURGE DEVICE command using the following syntax:

```
PURGE DEVICE device-name [/ALL]
```

To delete a device from the LAN volatile device database, enter the CLEAR DEVICE command using the following syntax:

```
CLEAR DEVICE device-name [/ALL]
```

For the PURGE DEVICE and CLEAR DEVICE commands, the /ALL qualifier deletes all LAN devices in the LAN permanent device database.

Examples

1.

```
LANCP> PURGE DEVICE/ALL
```

This command deletes all devices from the LAN permanent device database.

2.

```
LANCP> CLEAR DEVICE EXA0
```

This command deletes device EXA0 from the LAN volatile device database.

24.7 Managing the LAN Node Databases

The LAN volatile and permanent node databases contain a single entry for each defined LAN node. Each entry in the LAN volatile node database contains node information and MOP downline load counters information. Each entry in the LAN permanent node database contains node information that is used to populate the volatile database when the LANACP LAN server process is started.

Typically, each database contains the same nodes. The LANACP LAN server process maintains the volatile node database. The LANCP utility maintains the permanent node database. You can manipulate either database using the LANCP utility commands depending on your user privileges, as follows:

- Privileged users can add or delete node entries from each database and clear MOP downline load counters information for LAN nodes
- Unprivileged users can view the node information and MOP downline load status and counters information

The following sections describe how to enter nodes into and remove nodes from the LAN permanent and volatile node databases.

24.7.1 Displaying Nodes in the LAN Node Databases

To display information in the LAN permanent node database, enter the LIST NODE command using the following syntax:

```
LIST NODE node-name [/ALL]
```

To display information in the LAN volatile node database, enter the SHOW NODE command using the following syntax:

```
SHOW NODE node-name [/ALL]
```

For the LIST NODE and SHOW NODE commands, the /ALL qualifier displays data for all nodes in the LAN permanent or volatile node database.

24.7.2 Entering Nodes into the LAN Node Databases

To enter a node into the LAN permanent node database or to modify an existing entry, enter the DEFINE NODE command using the following syntax:

```
DEFINE NODE node-name [/qualifiers]
```

To enter a node into the LAN volatile node database or to modify an existing entry, enter the SET NODE command using the following syntax:

```
SET NODE node-name [/qualifiers]
```

Table 24–11 provides a brief description of the DEFINE NODE and SET NODE command qualifiers.

Managing the Local Area Network (LAN) Software

24.7 Managing the LAN Node Databases

Table 24–11 DEFINE NODE and SET NODE Command Qualifiers

Qualifier	Description
<i>/ADDRESS=node-address</i>	Associates a LAN address with the node name.
<i>/ALL</i>	Defines data for all nodes in the LAN permanent or volatile node database.
<i>/BOOT_TYPE=VAX_SATELLITE ALPHA_SATELLITE OTHER</i>	Indicates the type of processing required for downline load requests.
<i>/FILE=file-spec</i>	Supplies the file name you want to be provided when the downline load request does not include a file name.
<i>/PERMANENT_DATABASE</i> (SET command only)	Updates the node entries in the LAN volatile node database with any data currently set in the permanent database.
<i>/ROOT=directory-specification</i>	Supplies the directory specification to be associated with the file name.
<i>/SIZE=value</i>	Specifies the size in bytes of the file data portion of each downline load message.
<i>/V3</i>	Forces the server to respond to only MOP Version 3 boot requests from this node.
<i>/VOLATILE_DATABASE</i> (DEFINE command only)	Updates the node entries in the LAN permanent node database with any data currently set in the volatile database.

Examples

1.

```
DEFINE NODE GALAXY/ADDRESS=08-00-2B-11-22-33 -  
          /FILE=NISCS_LOAD.EXE -  
          /ROOT=$64$DIA14:<SYS10.> -  
          /BOOT_TYPE=VAX_SATELLITE
```

This command sets up node GALAXY in the LAN permanent node database for booting as a VAX satellite into an OpenVMS Cluster system.

The NISCS_LOAD.EXE file is actually located on \$64\$DIA14:<SYS10.SYSCOMMON.SYSLIB>. The <SYSCOMMON.SYSLIB> is supplied by the LANACP LAN server process and is not included in the root definition.

2.

```
DEFINE NODE ZAPNOT/ADDRESS=08-00-2B-11-22-33 -  
          /FILE=APB.EXE -  
          /ROOT=$64$DIA14:<SYS10.> -  
          /BOOT_TYPE=ALPHA_SATELLITE
```

This command sets up node ZAPNOT for booting as an Alpha satellite into an OpenVMS Cluster system.

The APB.EXE file is actually located on \$64\$DIA14:<SYS10.SYSCOMMON.SYSEXEXE>. Note that the <SYSCOMMON.SYSEXEXE> is supplied by the LANACP LAN server process and is not included in the root definition.

Managing the Local Area Network (LAN) Software

24.7 Managing the LAN Node Databases

3.

```
SET NODE CALPAL/ADDRESS=08-00-2B-11-22-33 -  
/FILE=APB_061.EXE
```

This command sets up node CALPAL for booting an InfoServer image. It defines the file that should be loaded when a load request without a file name is received from node CALPAL.

Because the file does not include a directory specification, the logical name LAN\$DLL defines where to locate the file. You could give a directory specification using the file name or by using the /ROOT qualifier.

Note that specifying the file name explicitly in the boot command overrides the file name specified in the node database entry.

24.7.3 Deleting Nodes from the LAN Node Databases

To delete a node from the LAN permanent node database, enter the PURGE NODE command using the following syntax:

```
PURGE NODE node-name [/ALL]
```

To delete a node from the LAN volatile node database, enter the CLEAR NODE command using the following syntax:

```
CLEAR NODE node-name [/ALL]
```

For the PURGE NODE and CLEAR NODE commands, the /ALL qualifier deletes all LAN nodes in the LAN permanent or volatile node database.

24.8 Understanding LAN MOP

The collection of utilities and startup command files for LANCP and LANACP provide the necessary functionality for MOP downline load service. These utilities and files load cluster satellites, terminal servers, and systems requiring downline load of special images, such as console update images or system software update images (for InfoServer load).

24.8.1 Coexistence with DECnet MOP

The LAN MOP environment provides functionality that is similar to that provided by DECnet. The result is that a system manager can choose which functionality to use, DECnet MOP or LAN MOP. For OpenVMS Cluster systems, LAN MOP permits the operation of a cluster without the presence of DECnet.

LAN MOP can coexist with DECnet MOP in the following ways:

- Running on different systems
For example, DECnet MOP service is enabled on some of the systems on the LAN, and LAN MOP is enabled on other systems.
- Running on different LAN devices on the same system
For example, DECnet MOP service is enabled on a subset of the available LAN devices on the system, and LAN MOP is enabled on the remainder.
- Running on the same LAN device on the same system but targeting a different set of nodes for service
For example, both DECnet MOP and LAN MOP are enabled, but LAN MOP has limited the nodes to which it will respond. This allows DECnet MOP to respond to the remainder.

Managing the Local Area Network (LAN) Software

24.8 Understanding LAN MOP

24.8.2 Migrating from DECnet MOP to LAN MOP

To migrate to LAN MOP, follow these steps:

1. Decide which nodes are to provide MOP downline load service. These may be the same nodes that currently have service enabled for DECnet.
2. Populate the LAN permanent device database by typing the following command at the DCL prompt:

```
MCR LANCP DEFINE DEVICE/UPDATE
```
3. Populate the LAN permanent node database by entering a node definition for each cluster satellite node and any other nodes that are similarly defined in the DECnet node database. You can enter this data manually or execute the command procedure `SYS$EXAMPLES:LAN$POPULATE.COM`, following the directions and help provided.
4. Disable service on each of the DECnet circuits where it is currently enabled in the volatile database.
5. Enable service on each LAN device in the LAN permanent device database that you would like to use by typing the following command at the DCL prompt for each device:

```
MCR LANCP DEFINE DEVICE device-name/MOPDLL=ENABLE
```
6. If high performance is required, select a data size of 1482 bytes and only reduce this if some load requests now fail. Alternatively, set up one system to load those clients that require a small data size and set up a different system to load the other clients.

To permanently migrate back to DECnet MOP, follow these steps:

1. Disable the MOP service in the volatile database by typing the following:

```
MCR LANCP SET DEVICE device-name/MOPDLL=DISABLE
```
2. Disable the MOP service in LANCP's permanent database by typing the following:

```
MCR LANCP DEFINE DEVICE device-name/MOPDLL=DISABLE
```
3. Reenable service on each DECnet circuit in the permanent and volatile databases.

Note

Any nodes that you added while booting with LAN MOP will not have been entered in the DECnet node database as targets for downline load, and they will need to be updated when you return to DECnet MOP.

24.8.3 Using `CLUSTER_CONFIG_LAN.COM` and LAN MOP

A cluster management command procedure has been provided to facilitate the use of LANCP for LAN MOP booting of satellites. Called `CLUSTER_CONFIG_LAN.COM`, it resides in `SYS$MANAGER` and is a direct parallel to `CLUSTER_CONFIG.COM`, which is used by cluster managers to configure and reconfigure an OpenVMS Cluster system. The two procedures perform the same functions, but `CLUSTER_CONFIG.COM` uses DECnet MOP for downline load, whereas `CLUSTER_CONFIG_LAN.COM` uses LAN MOP and does not use DECnet for anything. Therefore, when you add a new node, `CLUSTER_CONFIG_LAN.COM`

Managing the Local Area Network (LAN) Software

24.8 Understanding LAN MOP

does not ask for the node's DECnet node name and address. Instead, it queries for an SCS node name and an SCS node ID number.

For your convenience, you can still run CLUSTER_CONFIG.COM. When you execute CLUSTER_CONFIG.COM, it checks whether LANACP for MOP booting is also running. It also checks to see if DECnet is running. If LANACP is running and DECnet is not, then CLUSTER_CONFIG.COM dispatches to CLUSTER_CONFIG_LAN.COM. If CLUSTER_CONFIG.COM discovers that both LANACP and DECnet are running, it asks the user whether LAN MOP booting is being used, and whether it should call CLUSTER_CONFIG_LAN.COM for the user.

24.8.4 Sample Satellite Load

The following example shows how to issue commands to the LANCP utility to enable MOP downline load service and to define node ZAPNOT:

```
set acp/opcom
set device eza0/mopdll=enable
set node ZAPNOT/addr=08-00-2B-33-FB-F2/file=APB.EXE-
    /root=$64$DIA24:<SYS11.>/boot=Alpha
```

The following example shows the OPCOM messages displayed when you start up the LANACP LAN server process:

```
%%%%%%%%%% OPCOM 10-JAN-2001 06:47:35.18 %%%%%%%%%%%
Message from user SYSTEM on GALAXY
LANACP MOP Downline Load Service
Found LAN device EZA0, hardware address 08-00-2B-30-8D-1C

%%%%%%%%%% OPCOM 10-JAN-2001 06:47:35.25 %%%%%%%%%%%
Message from user SYSTEM on GALAXY
LANACP MOP Downline Load Service
Found LAN device EZB0, hardware address 08-00-2B-30-8D-1D

%%%%%%%%%% OPCOM 10-JAN-2001 06:47:54.80 %%%%%%%%%%%
Message from user SYSTEM on GALAXY
LANACP MOP V3 Downline Load Service
Volunteered to load request on EZA0 from ZAPNOT
Requested file: $64$DIA24:<SYS11.>[SYSCOMMON.SYSEXE]APB.EXE

%%%%%%%%%% OPCOM 10-JAN-2001 06:48:02.38 %%%%%%%%%%%
Message from user SYSTEM on GALAXY
LANACP MOP V3 Downline Load Service
Load succeeded for ZAPNOT on EZA0
System image, $64$DIA24:<SYS11.>[SYSCOMMON.SYSEXE]APB.EXE (Alpha image)
```

The following display shows the contents of the LAN\$ACP.LOG file:

```
10-JAN-2001 06:47:35.02 Found LAN device EZA0, hardware address
08-00-2B-30-8D-1C
10-JAN-2001 06:47:35.18 Found LAN device EZB0, hardware address
08-00-2B-30-8D-1D
10-JAN-2001 06:47:35.25 LANACP initialization complete
10-JAN-2001 06:47:45.39 Enabled LAN device EZA0 for MOP downline load service in
exclusive mode
10-JAN-2001 06:47:54.70 Volunteered to load request on EZA0 from ZAPNOT
Requested file: $64$DIA24:<SYS11.>[SYSCOMMON.SYSEXE]APB.EXE
10-JAN-2001 06:48:02.23 Load succeeded for ZAPNOT on EZA0
MOP V3 format, System image, $64$DIA24:<SYS11.>[SYSCOMMON.SYSEXE]APB.EXE
Packets: 2063 sent, 2063 received
Bytes: 519416 sent, 4126 received, 507038 loaded
Elapsed time: 00:00:07.42, 68276 bytes/second
```

Managing the Local Area Network (LAN) Software

24.8 Understanding LAN MOP

24.8.5 Cross-Architecture Booting

The LAN enhancements permit cross-architecture booting in a OPENVMS Cluster system. VAX boot nodes can provide boot service to Alpha satellites, and Alpha boot nodes can provide boot service to VAX satellites. Note that each architecture must include a system disk that is used for installations and upgrades.

24.9 Managing the LAN MOP Downline Load Service

The LANACP LAN server process maintains the LAN volatile node and device databases. The LANCP utility provides commands that:

- Display MOP downline load status and counters information
- Clear counters information
- Enable or disable OPCOM messages and packet tracing

Counters and status information is maintained for each node and device. Counters information includes transmitted and received byte and packet counts, transmit errors, logical errors such as protocol violations and timeouts, and number of load requests. Status includes the time of the last load and the status of the last load.

24.9.1 Enabling MOP Downline Load Service

To enable MOP downline load service, enter the SET DEVICE command using the following syntax:

```
SET DEVICE device-name/DLL=ENABLE
```

In this command, use the *device-name* parameter to supply the LAN controller device name.

See Section 24.6.2 for a complete description of this command.

24.9.2 Disabling MOP Downline Load Service

To disable MOP downline load service, enter the SET DEVICE command using the following syntax:

```
SET DEVICE device-name/DLL=DISABLE
```

In this command, use the *device-name* parameter to supply the LAN controller device name.

See Section 24.6.2 for a complete description of this command.

24.9.3 Displaying the Status and Counters Data

To display MOP downline load status, enter the SHOW DLL command using the following syntax:

```
SHOW DLL
```

The following display shows counters information for a particular node:

```
LAN MOP DLL Status:  
  EXA enabled in exclusive mode for known nodes only, data size 1482 bytes  
  FXA disabled
```

Managing the Local Area Network (LAN) Software

24.9 Managing the LAN MOP Downline Load Service

	#Loads	Packets	Bytes	Last load time	Last loaded
	-----	-----	-----	-----	-----
EXA	5	1675	4400620	10-JAN-2001 10:27.51	GALAXY
FXA	0	0	0		

On this node are two LAN devices, EXA (DEMNA) and FXA (DEMFA). MOP downline load service is enabled on EXA in exclusive mode.

Requests are answered only for nodes that are defined in the LANACP node database. The image data size in the load messages is 1482 bytes. There have been five downline loads, the last one occurring on node GALAXY at 10:27. Finally, no downline loads are recorded for FXA, which is currently disabled for downline load service.

To display recent downline load activity that has been logged in the LAN\$ACP.LOG file, enter the SHOW LOG command using the following syntax:

```
SHOW LOG
```

24.9.4 Displaying the Status and Counters Data for Individual Nodes

To display MOP downline load information for nodes in the LAN permanent node database, enter the LIST NODE command using the following syntax:

```
LIST NODE node-name [/qualifiers]
```

To display MOP downline load status and counters information for nodes in the LAN volatile node database, enter the SHOW NODE command using the following syntax:

```
SHOW NODE node-name [/qualifiers]
```

Table 24–12 provides a brief description of the LIST NODE and SHOW NODE command qualifiers.

Table 24–12 LIST NODE and SHOW NODE Command Qualifiers

Qualifier	Description
/ALL	Displays information for all nodes in the database.
/OUTPUT= <i>file-name</i>	Indicates that the output should be directed to the specified file. If the file name extension is .com, then the output is in the form of a list of DEFINE NODE or SET NODE commands. The resulting command file can be used to create the LAN node databases.
/TOTAL (SHOW NODE command only)	Displays counter totals only.

Example

The following example shows output from a command issued on a local node on which there are three nodes defined (GALAXY, ZAPNOT, and CALPAL). CALPAL has issued two load requests:

- The first request is the multicast request from CALPAL that the local node volunteered to accept.
- The second request is the load request sent directly to the local node by CALPAL for the actual load data. The elapsed time from the second load request to completion of the load was 6.65 seconds.

Managing the Local Area Network (LAN) Software

24.9 Managing the LAN MOP Downline Load Service

```
Node Listing:
GALAXY (08-00-2B-2C-51-28):
  MOP DLL: Load file:  APB.EXE
           Load root:  $64$DIA24:<SYS11.>
           Boot type:  Alpha satellite

ZAPNOT (08-00-2B-18-7E-33):
  MOP DLL: Load file:  NISCS LOAD.EXE
           Load root:  LAVC$SYSDEVICE:<SYS10.>
           Boot type:  VAX satellite

CALPAL (08-00-2B-08-9F-4C):
  MOP DLL: Load file:  READ_ADDR.SYS
           Last file:  LAN$DLL:APB_X5WN.SYS
           Boot type:  Other
           2 loads requested, 1 volunteered
           1 succeeded, 0 failed
           Last request was for a system image, in MOP V4 format
           Last load initiated 10-jan-2001 09:11:17 on EXA0 for 00:00:06.65
           527665 bytes, 4161 packets, 0 transmit failures

Unnamed (00-00-00-00-00-00):
Totals:
  Requests received      2
  Requests volunteered  1
  Successful loads       1
  Failed loads           0
  Packets sent           2080
  Packets received       2081
  Bytes sent             523481
  Bytes received         4184
  Last load              CALPAL at 10-jan-2001 09:11:17.29
```

24.9.5 Clearing the Counters Data

To clear MOP downline load counters for all nodes and devices, enter the CLEAR DLL command using the following syntax:

```
CLEAR DLL
```

24.9.6 OPCOM Messages

By default, OPCOM messages are enabled. Messages are generated by the LANACP LAN server process when device status changes, load requests are received, and loads complete. These messages are displayed on the operator's console and included in the log file written by LANACP, SYS\$MANAGER:LAN\$ACP.LOG.

To enable OPCOM messages, enter the SET ACP/OPCOM command using the following syntax:

```
SET ACP/OPCOM
```

24.9.7 Load Trace Facility

If the error data produced by the LANACP LAN server process for a load request is not sufficient to help you determine why the load is failing, you can direct the server process to record trace data. The data consists of transmit and receive packet information for every transmit and receive done by the server, and written to a log file for each load attempt. The name of the log file is SYS\$MANAGER:LAN\$nodename.LOG. You can record either all packet data or only the first 32 bytes of each packet.

Managing the Local Area Network (LAN) Software

24.9 Managing the LAN MOP Downline Load Service

The following list describes the typical load sequence:

1. Receive a Program Request message on the Load Assistance Multicast Address from the requesting node, code 8.
2. Transmit an Assistance Volunteer message to the requesting node, code 3.
3. Receive a Program Request message on your node address from the requesting node, code 8.
4. Transmit a Memory Load message to the requesting node with sequence number zero, code 2.
5. Receive a Request Memory Load message requesting the next sequence number (modulo 256), code 10 (decimal).
6. Repeat steps 4 and 5 until there is no more data to send.
7. Transmit a Memory or Parameter Load with Transfer Address message, code 0 or 20 (decimal).
8. Receive a final Request Memory Load message requesting the next sequence number (modulo 256) indicating that the last message has been received, code 10 (decimal).

For cluster satellite loads, the last Memory Load message contains cluster parameters. This message and the final Load with Transfer Address messages are displayed in full even if only partial trace echo has been enabled.

To enable partial tracing of packet data, enter the SET ACP/ECHO command using the following syntax:

```
SET ACP/ECHO
```

To enable full tracing of packet data, add the /FULL qualifier:

```
SET ACP/ECHO/FULL
```

24.9.8 MOP Console Carrier

Console carrier provides a mechanism to connect to a LAN device, such as a terminal server, that implements a management interface using the MOP console carrier protocol. The LANCP utility provides this function in the form of a CONNECT NODE command.

The command syntax is:

```
CONNECT NODE node-specification [/qualifiers]
```

Table 24–13 provides a brief description of the CONNECT NODE command qualifiers.

Table 24–13 CONNECT NODE Command Qualifiers

Qualifier	Description
/DEVICE= <i>device-name</i>	Specifies the LAN controller device name to be used for the connection.
/DISCONNECT= <i>disconnect-character</i>	Specifies a character that you can use to terminate the connection to the remote node.

(continued on next page)

Managing the Local Area Network (LAN) Software

24.9 Managing the LAN MOP Downline Load Service

Table 24–13 (Cont.) CONNECT NODE Command Qualifiers

Qualifier	Description
/PASSWORD=16hexdigits	Supplies the password to be used when the connection is initiated.
/V3 or /V4	Indicates that MOP Version 3 or Version 4 formatted messages, respectively, are to be used to make the connection.

Examples

1.

```
CONNECT NODE GALAXY/DEVICE=EWA0
```

This command attempts a console carrier connection to node GALAXY using the Ethernet device EWA0.

2.

```
CONNECT NODE 08-00-2B-11-22-33/DEVICE=EWA0/PASSWORD=0123456789ABCDEF
```

This command attempts a console carrier connection to the given node address using the Ethernet device EWA0, with a password.

24.9.9 MOP Trigger Boot

Some systems recognize and respond to MOP remote boot requests. These systems typically require a password or other mechanism to prevent unwanted boot requests from triggering a reboot of the system. The LANCP utility provides this function in the form of the TRIGGER NODE command.

To request a reboot of a LAN system, enter the TRIGGER NODE command using the following syntax:

```
TRIGGER NODE node-specification [/qualifiers]
```

Table 24–14 provides a brief description of the TRIGGER NODE command qualifiers.

Table 24–14 TRIGGER NODE Command Qualifiers

Qualifier	Description
/DEVICE= <i>device-name</i>	Specifies the LAN controller device name to be used for sending the boot messages.
/PASSWORD=16hexdigits	Supplies the password to be used when the connection is initiated.

Rather than specify the format to send MOP Version 3 or 4, the LANCP utility sends one message in each format to the target node.

The following examples show how to use the TRIGGER NODE command:

Examples

1.

```
TRIGGER NODE GALAXY/DEVICE=EWA0
```

This command sends MOP trigger boot messages to node GALAXY using Ethernet device EWA0.

Managing the Local Area Network (LAN) Software

24.9 Managing the LAN MOP Downline Load Service

2.

```
TRIGGER NODE 08-00-2B-11-22-33/DEVICE=EWA0/PASSWORD=0123456789ABCDEF
```

This command sends MOP trigger boot messages to the given node address using the Ethernet device EWA0, with indicated password.

Managing InfoServer Systems

This chapter describes InfoServer functions and InfoServer Client for OpenVMS software, which enables OpenVMS systems to access InfoServer device services. The chapter also describes the tasks you must perform to start the client software on your system and to make InfoServer devices available as public devices.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Establishing a server management session	Section 25.3
Starting InfoServer Client for OpenVMS software automatically	Section 25.5.3
Making InfoServer devices available automatically	Section 25.6.3

This chapter explains the following concepts:

Concept	Section
InfoServer functions	Section 25.1
LASTport protocols	Section 25.2
InfoServer Client for OpenVMS functions	Section 25.4
LASTCP utility functions	Section 25.5
LADCP utility functions	Section 25.6

25.1 Understanding InfoServer Functions

The InfoServer system is a high-performance **virtual device server**. It can make available, or **serve**, compact discs, read/write disks, magneto-optical (MO) devices, and tapes to client systems on the local area network (LAN). Systems running InfoServer Client software can connect to the virtual devices and use them as though they are locally attached devices.

Unlike a file server, the InfoServer system does not impose a file system on the virtual devices that it serves. For example, the InfoServer system can serve a disk with any type of on-disk file structure. The client system interprets the on-disk structure and uses its own native file system to access data. Multiple on-disk structures can be served by and accessed on a single InfoServer system at the same time.

Managing InfoServer Systems

25.1 Understanding InfoServer Functions

The InfoServer system can perform the following functions:

- Serve compact discs

The InfoServer system serves compact discs automatically, using a disc's volume label as the service name when the server is booted or when a disc is inserted into an InfoServer drive. You do not have to perform any management action. Client systems simply bind to and mount the disc under its volume label.

The InfoServer system can automatically serve to OpenVMS clients compact discs that are in ODS-2 format. High Sierra and ISO-9660 compact discs and other media types can be served manually through the InfoServer management interface.
- Serve Small Computer System Interface (SCSI) tapes

Using service names, the InfoServer system can serve SCSI tape devices to the network. Client systems can connect to these tape devices and use them as though they were locally attached devices.
- Serve read/write disk partitions

A **partition** is a logical subset of an InfoServer read/write disk. A single disk can be subdivided into several partitions, each of which can be served to the network independently. To remote client systems, these partitions appear to be whole disks. For example, a client system using InfoServer Client for OpenVMS software can access the partitions and use them as though they are local hard disks.
- Act as an initial load system for OpenVMS systems

The InfoServer system can downline load the primary bootstrap program to OpenVMS systems by responding to Maintenance Operation Protocol (MOP) requests. The server can locate MOP downline load files on the OpenVMS software distribution compact disc and copy them into temporary MOP partitions on an InfoServer-formatted read/write disk.

The initial system load (ISL) bootstrap program connects back to the software distribution compact disc and boots Standalone Backup. The Backup utility is then used to copy the OpenVMS operating system save sets from the compact disc to a read/write disk attached to the system. All subsequent OpenVMS boots are done from the local read/write disk.
- Downline load other products

You can use the InfoServer system to load any Ethernet product by file name; that is, the server does not require a Network Control Program (NCP) database entry to locate the requested file. For example, X terminal clients use the InfoServer system to downline load their system software. You can create a special MOP partition and copy the desired file to that partition. The server additionally supports downline loading of services by Ethernet address. Each InfoServer system can handle up to 100 simultaneous downline loads more efficiently than host-based downline loaders, which must start processes to assist in the load.

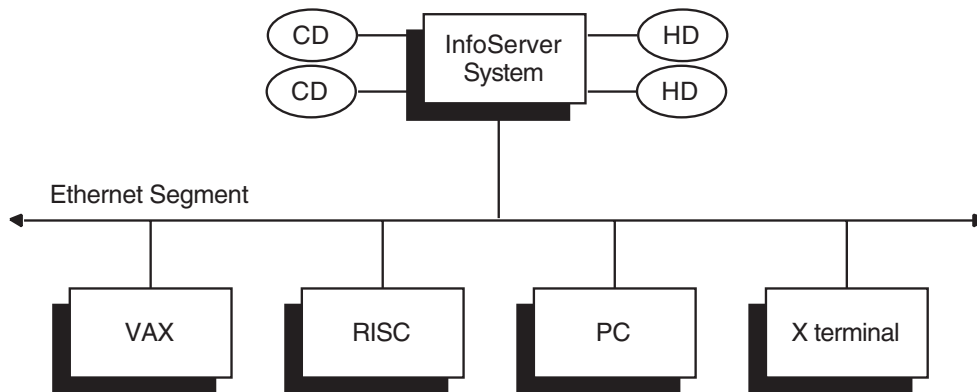
Figure 25–1 shows the relationship of the InfoServer system to several possible client systems. In this figure, two compact discs and two hard disks connected to the server appear to the client systems as local devices. The VAX system and the RISC workstation might be using one or two of the compact discs for software distribution and online documentation, while the PC might be referencing a disk

Managing InfoServer Systems

25.1 Understanding InfoServer Functions

partition on the InfoServer system. The X terminal boots from the InfoServer system and uses InfoServer disks for page, font, and customization files.

Figure 25–1 InfoServer System Serving Clients



ZK-5901A-GE

You can connect the InfoServer system to your Ethernet LAN and turn on the system. After the server is initialized, or **bootstrapped**, the server software automatically serves to client systems the device media connected to it. If you insert a compact disc into a server drive, the server detects this new device and automatically serves it to client systems by using the volume label as the service name.

The server bootstraps from its internal read/write device, on which the InfoServer software is preinstalled. InfoServer software updates are distributed on compact discs. As these new releases become available, you can install the software onto the internal device for subsequent booting. To update InfoServer software from the compact disc, follow these steps:

1. Insert the disc in a compact disc drive attached to the InfoServer system.
2. Move the InfoServer software to the internal read/write device. At the InfoServer prompt, enter a command in the following format, where *n* is the drive number:

On the InfoServer 100 or InfoServer 150 system:

```
InfoServer> UPDATE SYSTEM DKn:
```

On the InfoServer 1000 system:

```
InfoServer> UPDATE SYSTEM DKn: FLASH
```

The next time you boot the InfoServer system, it runs the updated software.

You can use the Software Products Library (formerly known as ConDIST) to update InfoServer software. After you log in to the InfoServer system, perform the following steps:

1. Insert the disk containing the [INFOSERVxxx] directory tree in a compact disk drive attached to the InfoServer system.

Managing InfoServer Systems

25.1 Understanding InfoServer Functions

2. At the InfoServer> prompt, enter a command in the following format, where *n* is the drive number:
 - On the InfoServer 100 or InfoServer 150 system, enter a command in the following format:
UPDATE SYSTEM DK*n*:
 - On the InfoServer 1000 system, enter a command in the following format:
UPDATE SYSTEM DK*n*: FLASH

These commands move the InfoServer software to the internal read/write device. The next time you boot the InfoServer system, it runs the updated software. Note that you can also boot the server from the Software Products Library disk.

You might want to customize server features. You can control InfoServer functions by logging in to the server and entering server commands, described in detail in the *InfoServer System Operations Guide*.

25.1.1 Automatic Service Policies for Multiple Servers

The InfoServer system automatically serves its locally connected devices to clients when the server is first powered on or when a removable device (for example, a compact disc) is inserted into a drive. The server reads the volume label of each device and uses the label as the name of the service offered to clients.

Note

You can disable the automatic service feature by using the InfoServer command SET SERVER AUTOMOUNT.

If multiple servers offer the same services, the client uses a rating scheme to select the appropriate service. Refer to the CREATE SERVICE command description in the *InfoServer System Operations Guide* for more information.

When you remove a compact disc from a server drive, the InfoServer system ends all client connections to the associated service. The InfoServer system also stops offering the associated service to client systems.

25.1.2 High-Availability Feature to Reduce Service Interruptions

The InfoServer system provides a high-availability feature that is especially beneficial for OpenVMS clients. If the server ends a service connection for some reason (for example, the server reboots, or you remove a compact disc), the OpenVMS client enters mount verification for that volume. If the same service is offered by another InfoServer system on the LAN, the client automatically connects to that service.

For example, suppose you have two identical copies of the OpenVMS Online Documentation compact disc in drives on two different servers. If one server or drive fails, a new connection is established to the duplicate disc on the other server. File operations continue as normal, and users experience almost no service disruption.

25.1.3 Support for X Terminal Clients

X terminal clients use the InfoServer system to download their system software, provide font services, save configuration information, and page memory to and from InfoServer disks. For example, system files for Compaq's VXT 2000 windowing terminals can be installed from compact disc on the InfoServer system. Once installed, these files are downline loaded on demand to each terminal when it is powered on.

The terminals can dynamically allocate partitions on an InfoServer disk as needed. For example, when a user requests that terminal customizations be saved, the InfoServer system automatically creates a disk partition to hold the information and creates a network service name for that partition. Once customization information is saved, the user can recall the information at any time.

VXT 2000 terminals that are InfoServer clients can also be virtual memory machines. Such terminals can page sections of main memory to and from InfoServer disks as required. Because a VXT client has no local disk, it uses InfoServer disks as page disks. When main memory is paged out to disk, the VXT client requests the InfoServer system to create a partition. This partition is then automatically extended as needed. Partitions and their network service names are created dynamically, without requiring user action.

By default, the InfoServer disk DK1, which is the internal disk that ships with each InfoServer 150 system, is enabled to allow VXT 2000 clients to allocate partitions remotely. Other disks can also be enabled through the use of InfoServer commands.

25.2 Understanding LASTport Protocols

The InfoServer system uses the LASTport transport protocol and the LASTport/Disk and LASTport/Tape system application protocols to provide access to the virtual devices it serves to the LAN. These protocols provide high-performance access to disk and tape devices. The InfoServer system implements the server portion of the protocols, while the client systems that access InfoServer storage devices implement the client portion.

On OpenVMS systems running the LASTport transport, all Ethernet devices must be terminated either by attaching the devices to an active network or by using an appropriate terminator. Failure to terminate the devices causes a system crash.

25.2.1 LASTport Transport Protocol

The **LASTport protocol** is a specialized LAN transport protocol that allows many clients to access InfoServer systems and perform reliable transactions. For the InfoServer system, a transaction is a device read or write operation. The LASTport protocol allows many client systems concurrently to read information from, and and write information to, an InfoServer storage device.

Unlike timer-based protocols, the LASTport protocol is a transaction-oriented protocol. Normally, information does not pass between a client and an InfoServer system unless the client initiates a transaction. The client system then runs a timer on the transaction, normally waiting from two to five seconds before assuming that the transaction is lost and retrying the operation.

Managing InfoServer Systems

25.2 Understanding LASTport Protocols

The LASTport protocol does not provide any routing functions; it runs only in a LAN. The LASTport protocol type is 80–41. If the extended LAN uses any filtering devices, the devices must allow this protocol type to pass unfiltered so that clients can access InfoServer systems across the filtering device.

The InfoServer system uses a multicast address feature of the LASTport protocol to establish connections to devices. The format of the multicast address is 09–00–2B–04–*nn–nn*, where *nn* depends on the work group enabled (refer to the *InfoServer System Operations Guide*).

25.2.2 LASTport/Disk Protocol

The LASTport/Disk protocol is a specialized device protocol that uses the LASTport transport. That is, LASTport/Disk messages are delivered in LASTport messages. The LASTport/Disk protocol provides the mechanism for reading and writing logical blocks independent from any underlying file system. The clients that implement the LASTport/Disk protocol interpret the file system locally. By using the LASTport/Disk protocol for access to compact discs and read/write disks, the InfoServer system can support multiple client operating systems and on-disk structures concurrently.

The LASTport/Disk protocol also provides the naming facility to access compact discs and read/write disks. The InfoServer system assigns each virtual device a service name and allows clients to query the LAN for these names. When the requested service is found, the client connects to it, and device access can begin. When duplicate virtual devices are available under identical service names, the protocol provides a facility for load balancing among the available devices.

25.2.3 LASTport/Tape Protocol

Like the LASTport/Disk protocol, the LASTport/Tape protocol uses the LASTport transport. That is, LASTport/Tape messages are delivered in LASTport messages. The LASTport/Tape protocol provides the mechanism for reading and writing tape records. Tape devices attached to the InfoServer system appear to tape clients as locally attached devices.

The LASTport/Tape protocol also provides the naming facility to access tapes. The InfoServer system assigns each tape device a service name and allows clients to query the LAN for these names. When the requested service is found, the client connects to it, and tape access can begin.

25.3 Establishing a Server Management Session

You can establish a server management session from a local or remote console terminal:

- **For a local session**, you connect a terminal capable of interpreting VT100 ANSI escape sequences to the serial port on the rear of the InfoServer system unit (MMJ1 on the InfoServer 150 unit). The terminal must be set to 9600 baud, 8 bits, no parity.
- **For a remote session**, you make a connection to the InfoServer system through a local area terminal (LAT) server.

Like many network servers, the InfoServer system advertises a LAT service for its management interface and accepts connections from remote terminals attached to terminal servers. Therefore, any terminal attached to a terminal server on the extended LAN can act as a console terminal for the InfoServer system (if the user knows the InfoServer management password).

Managing InfoServer Systems

25.3 Establishing a Server Management Session

Determining the Server's Default Service Name

To make a remote connection to the InfoServer system for the first time, you must determine the server's default name. To do this, add the four-character prefix LAD_ to the hexadecimal Ethernet datalink address on the InfoServer system's cabinet. You can change this default name by using the InfoServer command SET SERVER NAME.

The server's name is the LAT service name to which you connect. For example, if the default server name is LAD_08002B15009F, you would enter the following command at the terminal server's prompt to manage the InfoServer system:

```
Local> CONNECT LAD_08002B15009F
```

Refer to your terminal server user's guide for information about the establishment of LAT service connections.

Entering an InfoServer Password

After you connect to the InfoServer system, you must enter an InfoServer password to establish the management session. The default server password is ESS. You can change the password with the InfoServer command SET SERVER PASSWORD.

Example

The following example shows the establishment of a sample session using a DECserver 500 terminal server:

```
Local> CONNECT LAD_08002B133C1C
Password: ESS (not_ echoed)
Local -010- Session 1 to LAD_08002B133C1C established

DEC InfoServer V3.1
InfoServer> SHOW SERVER
```

In this example, the terminal server's prompt is Local>, and a LAT session is established to the InfoServer system whose service name is LAD_08002B133C1C. The InfoServer system prompts for a server password. When you enter the correct password, the server prompts for InfoServer commands with the InfoServer> prompt.

Ending a Session

At the end of the management session, you can enter the EXIT command at the InfoServer> prompt. This command returns you to the terminal server's Local> prompt if the management session is over a LAT connection.

25.3.1 Server Management Commands

Table 25-1 summarizes InfoServer commands and their functions.

Table 25-1 Summary of InfoServer Commands

Command	Function
BACKUP	Saves InfoServer-formatted disks.
BIND	Establishes a connection to the specified ODS-2 service and creates the virtual device VDK1 for that service.

(continued on next page)

Managing InfoServer Systems

25.3 Establishing a Server Management Session

Table 25–1 (Cont.) Summary of InfoServer Commands

Command	Function
CLEAR	Erases the console terminal screen.
COPY	Copies data from one disk or partition to another.
CRASH	Causes the server software to take a recognizable bugcheck, creating a dump if crashdump processing is enabled.
CREATE	Creates a new partition or service.
DELETE	Deletes a partition or service that was previously created.
DISCONNECT	Terminates a LASTport or LAT terminal server session.
ERASE	Erases the specified disk or partition; erases FUNCTIONS or SERVICES data from nonvolatile random-access memory (NVRAM).
EXIT	Terminates the management session.
HELP	Displays help text for the InfoServer commands.
INITIALIZE	Formats a read/write disk into an InfoServer disk.
LOOP	Automatically repeats any valid InfoServer command.
MONITOR	Automatically repeats valid InfoServer commands every 3 seconds, clearing the screen and placing the cursor at the home position.
PURGE	Purges old versions of VXT software.
REBOOT	Shuts down and reboots the server.
RECORD	Records data from an InfoServer disk or partition to a compact disc.
RESTORE	Resets the server to a previously saved system configuration.
RETRIEVE	Restores InfoServer-formatted disks saved by the BACKUP command.
REWIND	Rewinds an InfoServer tape.
SAVE	Saves configuration and service data for recovery after a server reboot.
SET	Sets partition, service, or server parameters.
SHOW	Displays the server's parameters and counters.
UNBIND	Deletes the VDK1 virtual device and terminates the connection to the remote service.
UNLOAD	Rewinds and unloads an InfoServer tape.
UPDATE	Installs one or more new software products or functions.
VERIFY	Validates the on-disk structure of a device formatted with the INITIALIZE command.
ZERO	Sets internal server counters to 0.

The InfoServer system provides a Help facility that contains information about each server command, including parameters, qualifiers, and examples of its use. For detailed information about InfoServer commands, refer to the *InfoServer System Operations Guide*.

25.4 Understanding InfoServer Client for OpenVMS Functions

InfoServer Client for OpenVMS software enables clients running the OpenVMS operating system to access virtual device services offered by InfoServer systems on a LAN. Software components include the following ones:

- **LASTport driver**
The LASTport driver provides reliable data transfer services for its clients. It interacts with the Data Link driver and the LASTport/Disk driver as an efficient transport for a virtual device service. The LASTport driver can support other applications, such as a primitive data queueing service.
- **LASTport/Disk client driver**
The LASTport/Disk client driver presents a standard block device interface to the system. The OpenVMS file system interacts with the LASTport/Disk client as if the LASTport/Disk client were a local disk driver. The LASTport/Disk client driver supports both raw and buffered interfaces.
- **LASTport/Tape client driver**
The LASTport/Tape client driver enables OpenVMS clients to access and use as local devices SCSI tapes attached to InfoServer systems.
- **LASTCP and LADCP utilities**
These utilities allow you to start InfoServer Client software on your system, monitor transport status, and configure and maintain InfoServer device services. Section 25.5 and Section 25.6 introduce the utilities. For complete information about the utilities, refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

25.5 Understanding LASTCP Utility Functions

InfoServer Client for OpenVMS software uses the LASTport protocol to communicate with InfoServer systems on an extended LAN. The protocol is implemented in the OpenVMS device driver ESS\$LASTDRIVER.

The LASTport Control Program (LASTCP) utility is the management interface that allows you to control and diagnose ESS\$LASTDRIVER. You can use LASTCP to perform the following tasks:

- Start and stop ESS\$LASTDRIVER
- Display counters for circuits, lines, nodes, and ESS\$LASTDRIVER
- Display node characteristics
- Display known clients and servers
- Display LASTport status
- Reset counters

The description of the LASTCP utility covers the following topics:

- Invoking and exiting the utility
- LASTCP command summary
- Starting InfoServer Client for OpenVMS software automatically

Managing InfoServer Systems

25.5 Understanding LASTCP Utility Functions

25.5.1 Invoking and Exiting the LASTCP Utility

Use of LASTCP requires normal privileges, except where noted. To invoke LASTCP, enter the following command:

```
$ RUN SYS$SYSTEM:ESS$LASTCP
%LASTCP-I-VERSION, ESS$LASTDRIVER V1.5 is running
LASTCP>
```

At the LASTCP> prompt, you can enter LASTCP commands. To exit the utility, enter EXIT or press Ctrl/Z at the LASTCP> prompt.

You can also execute a single LASTCP command by using a DCL string assignment statement, as shown in the following example:

```
$ LASTCP ::= $ESS$LASTCP
$ LASTCP SHOW CLIENTS
```

LASTCP executes the SHOW CLIENTS command and returns control to DCL command level.

25.5.2 LASTCP Command Summary

Table 25–2 summarizes LASTCP commands and their functions.

Table 25–2 Summary of LASTCP Commands

Command	Function
EXIT	Returns the user to DCL command level
HELP	Displays HELP text for LASTCP commands
SHOW CIRCUIT COUNTERS	Displays circuit counters
SHOW CLIENTS	Displays known clients
SHOW LINE COUNTERS	Displays line counters
SHOW NODE CHARACTERISTICS	Displays node characteristics
SHOW NODE COUNTERS	Displays node counters
SHOW SERVERS	Displays known servers
SHOW STATUS	Displays local status
SHOW TRANSPORT COUNTERS	Displays transport counters
START TRANSPORT	Starts LASTDRIVER
STOP TRANSPORT	Stops LASTDRIVER
ZERO COUNTERS	Resets counters

You can abbreviate LASTCP commands to the first unique characters of the command verb. For example, you can abbreviate the command SHOW SERVERS to SH SE.

LASTCP provides a Help facility that contains information about each command and its parameters and qualifiers, as well as examples of its use. For a complete description of LASTCP commands, refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

25.5.3 Starting InfoServer Client for OpenVMS Software Automatically

You must start InfoServer Client for OpenVMS software using the ESS\$STARTUP command procedure. To make sure the software is started automatically each time the system reboots, execute the startup procedure from within SYSTARTUP_VMS.COM.

How to Perform This Task

1. Determine the value of SCSNODE, your system's node name parameter. If the parameter is defined as the null string (the default value), InfoServer Client for OpenVMS software does not start.

If you are running or plan to run DECnet for OpenVMS, SCSNODE must be defined as the system's DECnet node name. If you do not plan to run DECnet, and if the system is an OpenVMS cluster member, SCSNODE must be defined as the SCS system name, a 1- to 8-character node name that is unique in the cluster.

To determine the value of SCSNODE, enter the following commands to invoke SYSMAN and display the parameter:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> PARAMETERS USE CURRENT
SYSMAN> PARAMETERS SHOW SCSNODE
```

2. If SCSNODE is defined as the null string, perform these steps:
 - a. Enter a command in the following format, where *node-name* is the system's DECnet node name or (if you do not plan to run DECnet for OpenVMS) the SCS system name:

```
PARAMETERS SET SCSNODE "node-name"
```

For example:

```
SYSMAN> PARAMETERS SET SCSNODE "MYNODE"
```

- b. Enter the following commands to write the new value to the parameter file and exit from SYSMAN:

```
SYSMAN> PARAMETERS WRITE CURRENT
SYSMAN> EXIT
```

- c. Add a line in the following format to the AUTOGEN parameter file SYS\$SYSTEM:MODPARAMS.DAT to define the SCSNODE parameter:

```
SCSNODE = "node-name"
```

For example:

```
SCSNODE = "MYNODE"
```

3. Invoke any editor to edit SYS\$MANAGER:SYSTARTUP_VMS.COM and find the command that starts InfoServer Client software. For example:

```
$ @SYS$STARTUP:ESS$STARTUP DISK
```

Note that the parameters CLIENT and DISK are synonymous. If the command is preceded by the DCL comment delimiter (!), remove the delimiter. To enable tape functions, add the TAPE parameter to the command line:

```
$ @SYS$STARTUP:ESS$STARTUP DISK TAPE
```

Managing InfoServer Systems

25.5 Understanding LASTCP Utility Functions

4. If SYSTARTUP_VMS.COM invokes the DECnet for OpenVMS startup procedure (SYS\$MANAGER:STARTNET.COM), make sure SYSTARTUP_VMS.COM executes the InfoServer Client for OpenVMS startup procedure *after* invoking STARTNET.COM.

The following example shows the network startup command line followed by the InfoServer Client for OpenVMS startup command line. Note that if you omit the TAPE parameter, only the disk function is started.

```
$ @SYS$MANAGER:STARTNET
.
.
.
$ @SYS$STARTUP:ESS$STARTUP DISK TAPE
```

5. Optionally, edit the file SYS\$STARTUP:ESS\$LAST_STARTUP.DAT to specify desired startup qualifiers for the LASTport transport. (Refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.)

25.5.4 InfoServer Client Can Fail to Start If DECnet Is Started or Stopped

The InfoServer client software fails to start on a system where DECnet has been started and subsequently stopped. The following message will be found in the file SYS\$MANAGER:ESS\$STARTUP.LOG:

```
%ESS-I-NONET ESS started before DECnet. 4-MAR-2000 16:36:39.29
```

If the InfoServer client must be started at this point, the LASTport transport can be started with the Last Control Program using the following command:

```
$ MCR ESS$LASTCP
LASTCP> START
```

This command will start the transport. You may now execute the InfoServer client startup:

```
$ @SYS$STARTUP:ESS$STARTUP DISK
```

Because the transport is already started, the startup will run successfully.

25.5.5 Multiple Controllers Configured But Not All Attached to Media (Alpha Only)

If you have multiple Ethernet and FDDI controllers configured on your OpenVMS Alpha system, you might experience problems with the InfoServer client transport (LASTport) under either of the following conditions:

- Not all of the Ethernet and FDDI controllers are connected to network cabling.
- An FDDI controller is connected to the network cabling, but the FDDI ring is not functional; for example, some FDDI hardware may be powered down or broken.

Problems can range from not being able to access all the services available on the network, if you have four or more controllers configured, to a system crash.

To avoid these problems, specify only the controllers that are attached to media. Compaq recommends that you do this by first editing your SYS\$STARTUP:ESS\$LAST_STARTUP.DAT data file to specify only the controllers that are attached and then restarting your system.

Managing InfoServer Systems

25.5 Understanding LASTCP Utility Functions

With certain controller configurations, if you specify controllers that are not attached, your system might crash when you issue the following command sequence:

```
$ MC ESS$LASTCP
LASTCP> STOP
```

An example of how to edit the SYS\$STARTUP:ESS\$LAST_STARTUP.DAT file follows. The unedited file is shown first, followed by an edited file.

```
!++
! This file will be used to set the appropriate LASTCP qualifiers. The following
! LASTCP qualifiers: ALL CONTROLLERS, CHECKSUM, TRANSMIT_QUOTA, or SLOW_MODE
! can be set by using the following statement format:
! LASTCP qualifier = 1 to enable e.g. SLOW_MODE = 1 enables SLOW_MODE
! LASTCP qualifier = 0 to disable e.g. SLOW_MODE = 0 disables SLOW_MODE
! The remaining LASTCP qualifiers will require the appropriate value settings.
! DEVICE = (list-of-devices)
! TIMEOUT = n minimum interval in seconds
! CIRCUIT_MAXIMUM = n maximum number of nodes
! GROUP = n Group number
! NODE_NAME = name Node name
! CONTROLLERS = ({{controller letter,}...}) Controller list
! TRANSMIT_QUOTA = n Number of transmit buffers
!--
ALL_CONTROLLERS = ON
```

The edited SYS\$STARTUP:ESS\$LAST_STARTUP.DAT file follows. This example assumes you have ESA, ETA, EXA, EZA controllers configured on your system and that only the ESA controller is attached to the Ethernet wire.

```
!++
! This file will be used to set the appropriate LASTCP qualifiers. The following
! LASTCP qualifiers: ALL CONTROLLERS, CHECKSUM, TRANSMIT_QUOTA, or SLOW_MODE
! can be set by using the following statement format:
! LASTCP qualifier = 1 to enable e.g. SLOW_MODE = 1 enables SLOW_MODE
! LASTCP qualifier = 0 to disable e.g. SLOW_MODE = 0 disables SLOW_MODE
! The remaining LASTCP qualifiers will require the appropriate value settings.
! DEVICE = (list-of-devices)
! TIMEOUT = n minimum interval in seconds
! CIRCUIT_MAXIMUM = n maximum number of nodes
! GROUP = n Group number
! NODE_NAME = name Node name
! CONTROLLERS = ({{controller letter,}...}) Controller list
! TRANSMIT_QUOTA = n Number of transmit buffers
!--
ALL_CONTROLLERS = OFF
DEVICE = (ESA)
```

Note

The default ESS\$LAST_STARTUP.DAT file is stored in SYS\$COMMON:[SYS\$STARTUP]. You might want to put the edited file in SYS\$SPECIFIC:[SYS\$STARTUP]. Otherwise, other system roots might be affected.

Managing InfoServer Systems

25.5 Understanding LASTCP Utility Functions

25.5.6 Startup Restrictions: PATHWORKS and RSM

If PATHWORKS or Remote System Manager (RSM) or both are installed, the InfoServer Client for OpenVMS startup must be run before the startup for PATHWORKS or RSM, or both. For example:

```
$ @SYS$MANAGER:STARTNET
.
.
.
$ @SYS$STARTUP:ESS$STARTUP DISK TAPE
$ @SYS$STARTUP:PCFS_STARTUP
$ @SYS$STARTUP:RSM$SERVER_STARTUP
```

InfoServer Client for OpenVMS software provides device drivers and control programs that are shared by both the PATHWORKS and RSM products. All InfoServer Client for OpenVMS components are prefixed with ESS\$. The drivers and control programs supplied with InfoServer Client for OpenVMS software provide all necessary support for both PATHWORKS and RSM in addition to InfoServer Client support. You must execute the InfoServer Client for OpenVMS startup in the site-specific startup before executing either the PATHWORKS or RSM startup procedure.

25.5.7 Startup Restrictions: SYSMAN

You cannot start InfoServer Client for OpenVMS from a subprocess. Because the OpenVMS System Management utility (SYSMAN) uses subprocesses to complete its tasks on remote nodes, SYSMAN cannot be used to execute the SYS\$STARTUP:ESS\$STARTUP procedure.

25.5.8 User Account Requirements

To work with InfoServer Client for OpenVMS software, user accounts on your system must have the following privileges and quotas:

- GRPNAM privilege to use the /GROUP qualifier of the LADCP command BIND; SYSNAM privilege to use the command's /SYSTEM qualifier.
- At a minimum, default UAF account quotas.

Refer to the AUTHORIZE section in the *OpenVMS System Management Utilities Reference Manual* for an explanation of how to verify and change account privileges and quotas.

25.5.9 System Parameter MAXBUF Requirement

To use all the LASTport Control Program (LASTCP) utility's SHOW functions, you must set the value of the system parameter MAXBUF to 32000 or greater.

25.6 Understanding LADCP Utility Functions

Use the LAD Control Program (LADCP) utility to configure and control the LASTport/Disk and LASTport/Tape protocols on OpenVMS systems. OpenVMS systems that use LASTport/Disk and LASTport/Tape services are called client systems. You can use LADCP to perform the following tasks:

- Establish **bindings** to services. A binding creates a new DAD n : virtual disk unit or a new MAD n : virtual tape unit on the local OpenVMS system.
- Remove bindings to services.

Managing InfoServer Systems

25.6 Understanding LADCP Utility Functions

You can control service access by using a service access password. You can also write-protect services. In this case, local OpenVMS users of a DAD*n*: or MAD*n*: device unit receive an error if they attempt a write operation to the unit.

The protocols allow you to access storage devices that reside on an InfoServer system as though they are locally connected to your OpenVMS system. Thus, several OpenVMS client systems can share the same read-only media, eliminating the need for duplicate drives and media.

DAD*n*: and MAD*n*: device units are also referred to as **virtual device units**. They represent the local OpenVMS context for a volume that resides on a remote server. The OpenVMS driver that controls the DAD*n*: units is called ESS\$DADDRIVER. The OpenVMS driver that controls the MAD*n*: units is called ESS\$MADDRIVER.

The LASTport/Disk and LASTport/Tape protocols depend on the LASTport transport. The ESS\$STARTUP.COM command procedure in SYS\$STARTUP automatically loads ESS\$DADDRIVER and ESS\$MADDRIVER as well as ESS\$LASTDRIVER, the LASTport transport driver.

Note

Your site-specific startup command procedure must include a call to ESS\$STARTUP.COM. If you are using DECnet software, you must place the call *after* the @SYS\$MANAGER:STARTNET.COM command that starts DECnet software. See Section 25.5.3.

25.6.1 Invoking and Exiting the LADCP Utility

To invoke LADCP, enter the following command:

```
$ RUN SYS$SYSTEM:ESS$LADCP
LADCP>
```

You can enter LADCP commands at the LADCP> prompt.

You can also execute a single LADCP command by using a DCL string assignment statement, as shown in the following example:

```
$ LADCP ::= $ESS$LADCP
$ LADCP BIND CD_DOC_00661 /NOWRITE
```

LADCP executes the BIND command and returns control to DCL command level.

To exit LADCP, enter EXIT or press Ctrl/Z after the LADCP> prompt.

25.6.2 LADCP Command Summary

Table 25–3 summarizes LADCP commands and their functions.

Table 25–3 Summary of LADCP Commands

Command	Function
BIND	Establishes a service binding and creates a device unit
DEALLOCATE	Terminates any active connection to a service without deleting the unit control block (UCB)

(continued on next page)

Managing InfoServer Systems

25.6 Understanding LADCP Utility Functions

Table 25–3 (Cont.) Summary of LADCP Commands

Command	Function
EXIT	Returns the user to DCL command level
HELP	Displays help text for LADCP commands
SHOW SERVICES	Displays services offered by InfoServer systems on the LAN
UNBIND	Terminates an established service binding

LADCP provides a Help facility that contains information about each LADCP command, including parameters, qualifiers, and examples of its use. For detailed descriptions of LADCP commands, refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

25.6.3 Making InfoServer Devices Available Automatically

You can make remote InfoServer devices available on your system each time the system boots. To do so, add to SYSTARTUP_VMS.COM a series of LADCP BIND commands. For more information about the BIND command, refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

How to Perform This Task

1. Edit SYSTARTUP_VMS.COM and find the command that starts InfoServer Client software. For example:

```
@SYS$STARTUP:ESS$STARTUP DISK TAPE
```

This command starts the software with disk and tape functions.

2. Add the following command to invoke LADCP:

```
$ RUN SYS$SYSTEM:ESS$LADCP
```

3. Immediately after this command, add BIND commands in the following format to make InfoServer compact discs or read/write disks available as virtual device units:

```
BIND [/QUALIFIER,...] service-name
```

To make tape devices available, you must specify the /TAPE qualifier in addition to any other desired qualifiers:

```
BIND/TAPE [/QUALIFIER,...] service-name
```

For *service-name*, specify the name of the InfoServer device service. Usually a service name is the label of the volume to which the InfoServer system is providing access. For more information about the BIND command, refer to the *InfoServer Client for OpenVMS LASTCP and LADCP Utilities* manual.

4. Add an EXIT command to exit LADCP.
5. Add MOUNT commands in the following format to make available as public devices the virtual device units created in the previous step:

```
MOUNT/SYSTEM/NOASSIST device-name volume-label
```

For *device-name*, specify the name of the device. For *volume-label*, specify a volume label to assign to the device. For more information about the MOUNT command, refer to the MOUNT section in the *OpenVMS System Management Utilities Reference Manual*.

Managing InfoServer Systems

25.6 Understanding LADCP Utility Functions

Example

The following commands, executed in SYSTARTUP_VMS.COM, start the InfoServer Client software and make available to client systems the InfoServer device DAD\$OPENVMSV72.

```
.  
. .  
$ @SYS$STARTUP:ESS$STARTUP DISK  
$ RUN SYS$SYSTEM:ESS$LADCP  
  BIND OPENVMSV72  
  EXIT  
$ MOUNT/SYSTEM/NOASSIST DAD$VMS055 VMS055  
. .  
.
```

In this example, the OpenVMS Version 7.2 consolidated distribution (CONDIST) compact disc loaded in a compact disc drive connected to an InfoServer system, is made available on the server as a virtual device unit and mounted as a public device.

Managing the LAT Software

This chapter describes how the LAT software works and the tasks you must perform to implement and manage the LAT software on your system.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Starting up the LAT protocol	Section 26.5
Customizing LAT characteristics	Section 26.6
Creating a service	Section 26.6.1
Setting up ports	Section 26.6.2
Setting up printers	Section 26.6.2.1
Setting up special application services	Section 26.6.2.2
Enabling queued incoming requests	Section 26.6.3
Enabling outgoing LAT connections	Section 26.6.4
Managing the LATACP database size	Section 26.7

This chapter explains the following concepts:

Concept	Section
LAT protocol	Section 26.1
LAT network	Section 26.2
LAT configurations	Section 26.3
LAT Control Program utility	Section 26.4

26.1 Understanding the LAT Protocol

The operating system uses the LAT (local area transport) software to communicate with **terminal servers** and other systems within a local area network (LAN). Terminal servers are communication devices dedicated for connecting terminals, modems, or printers to a LAN. They offer the following features:

- Provide a cost-effective method of connecting many user terminals to a computer
- Save on cable requirements
- Maximize the number of devices that can access a computer

Managing the LAT Software

26.1 Understanding the LAT Protocol

With the LAT software, which implements the LAT **protocol**, the operating system can offer resources, or **services**, that the terminal servers can access. A system that offers LAT services is called a **service node**. In addition, nodes can access LAT services by enabling outgoing connections (using LATCP) and using the DCL command SET HOST/LAT. (In the remainder of this chapter, “servers” refers both to dedicated terminal servers and to nodes that allow outgoing access to other LAT services.)

A LAT service can consist of all the resources of a computer system, or it can be a specific resource on a computer system, such as an application program. You can set up your system as a **general timesharing service**, meaning that all of its resources are available to users in the LAN, or you can restrict access to a specific service (application program) on the system. This chapter and the *OpenVMS I/O User's Reference Manual* outline the procedure you use to set up access to a dedicated application program.

26.1.1 How the LAT Protocol Works

The LAT protocol allows the terminal servers and computers to communicate within a LAN, such as the Ethernet or the Fiber Distributed Data Interconnect (FDDI). The LAT protocol matches terminals and other devices to the computing resources (services) of the LAN. Because LAT terminals are not connected directly to the computer (service node) they are accessing, the local server must listen for service requests from its terminals and be able to match the terminals with computers that provide the desired services.

Using the LAT protocol, then, the operating system announces its available services over the LAN. Servers listen to the LAN announcements and build a database of service information so that they can locate an appropriate system when a user terminal requests computing services. For example, a user terminal might request general processing service or a data entry program on the operating system. A server uses the LAT protocol to establish and maintain a connection between the requesting terminal and the operating system.

Sometimes the operating system can request services from a terminal server. The LAT protocol allows systems to ask for connections to printers or other devices attached to a terminal server.

26.1.2 Advantages of the LAT Protocol

Using the LAT protocol on your system has many advantages:

- The LAT protocol lets you make the resources of any computer on a local area network available to any user in that network.
- In addition to general processing resources, you can set up terminals, printers, and modems so they are available from multiple systems in the local area network. This lets you efficiently use these resources and keep them available even if one of the systems in the network must be shut down.
- You can also set up application programs, such as data entry programs or news services, as resources. When a user requests a connection to the resource, the LAT protocol sets up a connection directly to the application program. No login procedure is necessary.
- The LAT protocol provides load balancing features and recovery mechanisms so users get the best, most consistent service possible. In their broadcast messages, systems rate the availability of their services so that servers can establish connections to computing resources on the least busy node. If a node

becomes unavailable for any reason, the servers attempt to provide access to alternate services.

- Users can establish multiple computing sessions on their terminals, connecting to several different computers and switching easily from one computing session to another. After switching from one session to another, users can return to the previous session and pick up where they left off. This saves users the time normally required to close out and reopen files or accounts and to return to the same point in a session.
- Finally, the LAT protocol can provide improved system performance. Because the servers bundle messages onto a single LAN interface, a server interface decreases the network traffic and reduces the number of computer interrupts realized in systems where terminals, modems, and printers each have a physical connection to the computer.

26.2 Understanding the LAT Network

A **LAT network** is any local area network where terminal servers and operating systems use the LAT protocol. A LAT network can coexist on the same LAN with other protocols. The LAT protocol, which operates on both terminal servers and the operating systems, is designed to ensure the safe transmission of data over the LAN.

The LAT network consists of the following components:

Component	For More Information
Service nodes	Section 26.2.1
Terminal server nodes	Section 26.2.2
Nodes allowing outgoing connections	Section 26.2.3
LAN cable	Section 26.2.4

Service nodes supply computing resources for the local network, while terminal server nodes (or nodes allowing outgoing connections) port their terminals, modems, or printers to those resources upon request from a user terminal or an application program.

Note that in a LAT network, nodes that *access* services are often referred to as *master* nodes, which distinguishes them from nodes that only *provide* services.

You can use the LAT Control Program (LATCP) to configure the LAT characteristics for your system. LATCP allows you to set up your system to support:

- Incoming access only
- Outgoing access only
- Both incoming and outgoing access

The systems that support incoming LAT connections are **service nodes**. (Using LATCP, you can also set up your system so that it supports neither incoming nor outgoing access.)

Managing the LAT Software

26.2 Understanding the LAT Network

26.2.1 Service Nodes

A service node is one type of node in a LAT network. (Nodes that are not running an OpenVMS operating system can also be used along with the OpenVMS nodes in a LAT network.) A service node is an individual computer in a LAN that offers its resources to users and devices. Because the OpenVMS operating systems contain the LAT protocol, any OpenVMS system can be configured as a service node within a LAT network.

26.2.1.1 Types of Services

Each node offers its resources as a **service**. Often, a node offers a general processing service, but it can offer limited services or special application services as well. Any or all of the services can be specialized applications.

For example, your service node might offer services for the following items:

- General processing
- Data entry
- Stock quotations

The general processing service would allow the use of the general computing environment. The data entry and stock services, on the other hand, would be restricted environments, with connections to the application service but to no other part of the service node.

Each service is distinguished by the name the system manager assigns to it. In an OpenVMS Cluster, Compaq recommends that the service name be the same as the cluster name. In an independent node, Compaq recommends that the service name be the same as the node name. With special service applications, the service holds the name of the application.

26.2.1.2 Service Announcements

A service node announces its services over the LAN at regular intervals so that terminal servers (and OpenVMS systems that allow outgoing connections) know about the availability of these network resources. The service announcement provides the physical node name, the service names, a description of services, and a rating of service availability. Servers listen to the LAN announcements and record information in a database. On nodes allowing outgoing connections, this database is maintained by the LAT Ancillary Control Process (LATACP). (See Section 26.7 for more information about managing the LATACP database.)

Whenever a user terminal or application program requests a service, the server node connects to the appropriate service node.

Note that you can disable a local node from multicasting service announcements by using the `/NOANNOUNCEMENTS` qualifier with the `LATCP` command `SET NODE`. However, because remote nodes must rely on the LAT service responder feature in the LAT protocol Version 5.2 (or higher) to connect to the local node, Compaq recommends that you use this qualifier only in a networking environment where newer model terminal servers and hosts are present (all LAT hosts, terminal servers, and PCs are running at least Version 5.2 of the LAT protocol). Otherwise, systems running versions of the LAT protocol prior to Version 5.2 (for example, DECserver 100, 200, and 500 systems) will be unable to connect to any of the systems that have LAT service announcements disabled.)

26.2.1.3 Print Requests

In some cases, service nodes can request services from terminal servers. The most common situation is when the system wants to use a printer that is connected to a terminal server port. The system submits the print request to the terminal server print queue that is set up and initialized in the OpenVMS startup procedure. Then the LAT symbiont (the process that transfers data to or from mass storage devices) requests the LAT port driver to create and terminate connections to the remote printer.

For information about setting up queues for printers connected to LAT ports, see Section 14.1.3 and Section 14.2.2.4.

26.2.2 Terminal Server Nodes

A **terminal server node** is the second type of node in a LAT network. A terminal server node is usually located near the terminals and printers it supports. The terminals and printers are physically cabled to the terminal server; the terminal server is physically connected to the LAN cable.

26.2.2.1 Locating Service Nodes

Terminal servers build and maintain a directory of services from announcements advertised over the network. Then, when terminal servers receive requests from terminal users, they can scan their service databases and locate the computer that offers the requested service.

Terminal servers not only look for the node that provides the requested service, but they can also evaluate the service rating of that node. If a requested service is offered by more than one node, then the service rating is used to select the node that is least busy. A server establishes a logical connection between the user terminal and the service node.

26.2.2.2 Setting Up Connections

One logical connection carries all the data directed from one terminal server node to a service node. That is, the server combines data from all terminals communicating with the same node onto one connection. A terminal server establishes a logical connection with a service node only if a logical connection does not already exist.

If a connection fails for any reason, a terminal server attempts to find another node offering the same service and “rolls over” the connection so users can continue their computing sessions.

Even though terminal connections are bundled together, each terminal can be uniquely identified by its name. A terminal name consists of two parts: the first part is the name of the port on the terminal server that the terminal line plugs into; the second part is the name of the terminal server node.

26.2.2.3 Servicing Nodes

Although terminal servers are usually the requesting nodes in a LAT network, sometimes service nodes request service from terminal servers. Most commonly, a service node queues print requests to remote printers connected to terminal servers.

Managing the LAT Software

26.2 Understanding the LAT Network

26.2.3 Nodes Allowing Outgoing Connections

Nodes can be set up to allow incoming connections, outgoing connections, or both. Nodes (excluding those that offer incoming connections only) such as terminal server nodes can locate service nodes and set up connections. The database of information about available nodes and services is maintained by the LAT Ancillary Control Process (LATACP). (See Section 26.7 for more information about managing the LATACP database.)

On a node that is set up to allow outgoing LAT connections, a user can connect to another node in the LAT network by entering the SET HOST/LAT command. For more information, refer to the SET HOST/LAT command in the *OpenVMS DCL Dictionary*.

26.2.4 Components of a LAT Network

Figure 26–1 is an example of a LAT network. The network consists of an Ethernet cable connecting service nodes and terminal server nodes.

The three service nodes in Figure 26–1, named MOE, LARRY, and ALEXIS, each offer services to terminal server nodes on the network.

Two of the service nodes, MOE and LARRY, belong to the OFFICE cluster. (The cluster is distinguished by its computer interconnect [CI] and star coupler.) Because MOE and LARRY are clustered, their service names are the same as their cluster name. Because both service nodes offer an OFFICE service, terminal server nodes can assess the work load on both OFFICE nodes and establish a connection to a node that offers the service that is less busy.

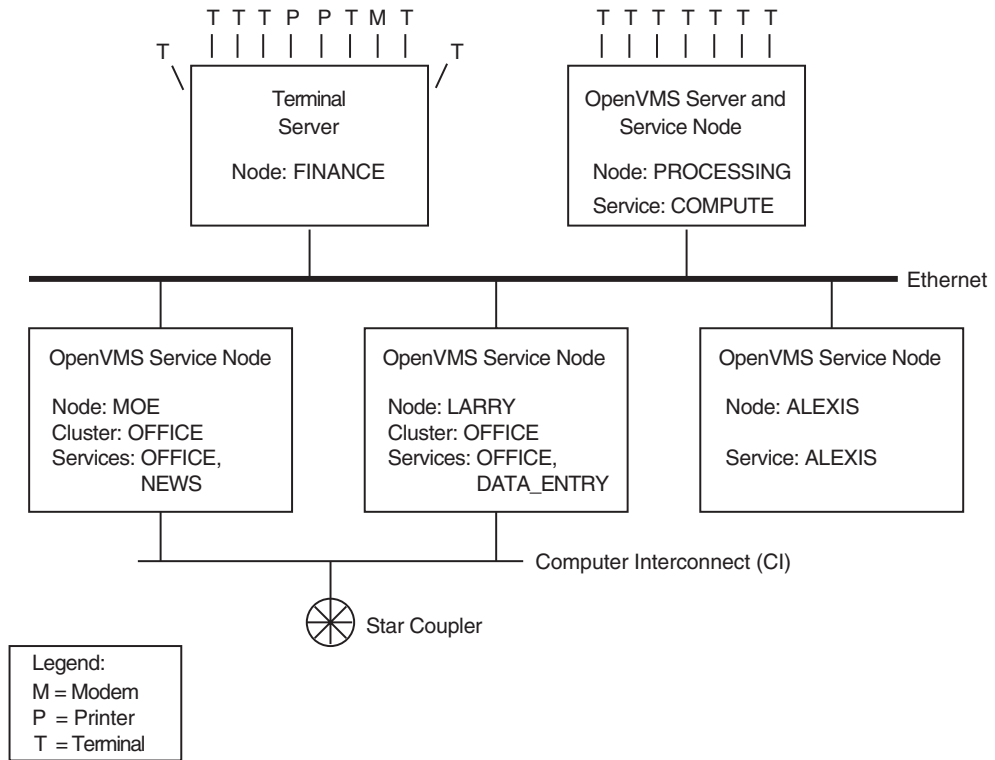
The third service node, ALEXIS, is an independent node in the LAT network so its service name is the same as its node name.

In addition to its primary OFFICE service, node MOE offers an application service called NEWS. With this specialized service, user terminals can connect directly to the online news program, without any login procedure but also without general access to the general computer resources of the node.

The node FINANCE, shown in Figure 26–1, is a terminal server node; it supports a number of interactive terminals, a modem, and printers. The node PROCESSING allows outgoing connections; it supports several interactive terminals. The node FINANCE can accept print requests from any of the three service nodes, provided each service node has set up print queues to support remote printers on the terminal server.

Node PROCESSING is also a service node. It offers the service COMPUTE.

Figure 26–1 Components of a LAT Network



ZK-1110A-GE

26.3 Understanding LAT Configurations

When you set up a LAT system, you need to understand the relationship between the LAT software and the network so you can configure your system to operate efficiently. The following sections contain information that will help you understand the following concepts:

- The relationship between LAT software and OpenVMS Cluster software
- The relationship between LAT software and DECnet software
- How the LAT software works in a networking environment that uses multiple LAN adapters
- How to use the LAT software in an Ethernet/FDDI configuration that cannot use large buffers

26.3.1 LAT Relationship to OpenVMS Clusters and DECnet

Although the LAT protocol works independently of OpenVMS Cluster software, Compaq recommends that you configure a service node to complement the OpenVMS Cluster concept. You achieve this by creating a service on each node in an OpenVMS Cluster and assigning the cluster name to this service. A terminal server assesses the availability of cluster services and establishes a connection to the node that is least busy. Thus, the LAT protocol helps balance the cluster load. If one node in the cluster fails, the terminal server can transfer the failed connections to another service node within the cluster.

Managing the LAT Software

26.3 Understanding LAT Configurations

The LAT software does not use DECnet as a message transport facility, but instead uses its own virtual circuit layer to implement a transport mechanism. The LAT and DECnet software work independently in a common LAN environment. For compatibility, if a service node is also a DECnet node, the service node name should be the same as the DECnet node name.

26.3.1.1 LAT and DECnet Running on the Same Controller

If Ethernet ports will be running both DECnet and LAT, you must start the DECnet software *before* the LAT software. If you do not start DECnet software first, all existing LAT connections may terminate, and reconnecting to the system via LAT may not be possible.

26.3.1.2 LAT and DECnet Running on Different Controllers

If DECnet is configured on the system (or if the system is part of a cluster), the SCSSYSTEMID system parameter may contain a nonzero value. Normally, this is not a problem unless the system has two or more LAN controllers connected to the same logical LAN.

For example, if your system has an FDDI controller and an Ethernet controller, your site may be configured so that the FDDI ring attached to the FDDI controller and the Ethernet segment attached to the Ethernet controller are bridged by a 10/100 LAN bridge (FDDI-to-Ethernet). In this configuration, it is impossible to run LAT over both controllers.

In such a configuration, you *must* run LAT and DECnet over the same controller if SCSSYSTEMID is not 0. If they do not run on the same controller, DECnet starts first, which in turn causes the LAT startup on the other controller to fail. This failure occurs because LAT startup tries to use the AA-00-04-00-xx-xx address (the DECnet LAN address); however, because DECnet is already using this address on another controller, the data link layer prevents the LAT startup from using that address. (In a single logical LAN, all data link addresses must be unique. Because both controllers try to use the same address, it is no longer unique.)

Using the following command to create the LAT link also fails because the LAN driver tries to use the address based on SCSSYSTEMID:

```
LATCP> CREATE LINK LAT$LINK_2 /NODECNET
```

If SCSSYSTEMID is set to 0, configuring LAT and DECnet on different controllers is possible. However, in a cluster environment, SCSSYSTEMID cannot be set to 0.

26.3.2 Using Multiple LAN Adapters

When you use multiple LAN addresses for one LAT node, you can configure a system with multiple LAN adapters connected to the same logical LAN. The LAT software can run over each adapter simultaneously and can better maintain connections. For example, when a virtual circuit chooses a primary path and uses it for all LAT message transmissions, the LAT software can continue communications through another adapter or logical path if that original path becomes blocked.

Note


Nodes running versions of LAT software prior to Version 5.3 of the LAT protocol (included in the OpenVMS operating system beginning with Version 7.0) may exhibit some differences in behavior. Therefore, if your configuration includes earlier versions of the LAT software, such as Version 5.1 or Version 5.2, note the differences and considerations discussed in this chapter.

26.3.2.1 Supported Configurations

Although it is possible to run LAT over multiple LAN adapters, it is still not possible to route LAT from one logical LAN to another. The following examples show supported LAT configurations for nodes running Version 5.3 of the LAT protocol (including nodes running Version 5.2 and 5.1 as well).

The widely used configuration shown in Figure 26–2 has an OpenVMS system running LAT Version 5.3 software over two Ethernet adapters (labeled A and B in the diagram) connected to the same physical LAN as a DECserver 200.

Figure 26–2 Multiple Address LAT Configuration: One LAN with Mixed-Version LAT Nodes



When a LAT connection is started between the DECserver 200 and the OpenVMS system, the LAT software determines that it is possible to use both adapters A and B for the LAT virtual circuit. One of the adapters will be chosen as the primary communications path while the other will be present in the unlikely event that the primary path fails.

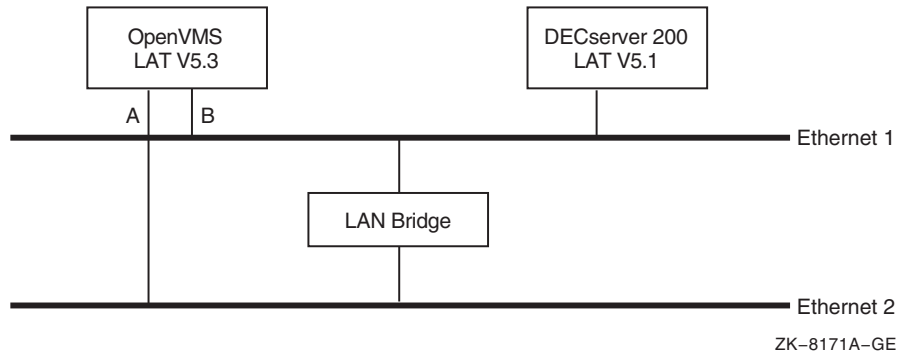
For example, if a user connects to the OpenVMS system from the DECserver 200, the OpenVMS system determines that there are two paths but chooses adapter B as the primary communications path. If the user runs a program that generates a large amount of output from the OpenVMS system and adapter B fails in some manner during the output, the LAT software will attempt to continue communications from the OpenVMS system to the DECserver through adapter A.

Figure 26–3 shows two LANs bridged together. However, this configuration will have the same characteristics as the configuration shown in Figure 26–2.

Managing the LAT Software

26.3 Understanding LAT Configurations

Figure 26–3 Multiple Address LAT Configuration: Two LANs with Mixed Version LAT Nodes




Note

It is possible for Ethernet 2 in Figure 26–3 to be an FDDI network. The LAT software regards each adapter as a network path with equal point-to-point communications and does not treat FDDI controllers any differently. However, for large buffer support, see Section 26.3.3 for more details.

In Figure 26–4, any virtual circuit created between the two OpenVMS systems will have two paths: through controllers B and C or A and D. If one path fails, the virtual circuit will continue over the other path. If both paths fail, the virtual circuit will eventually time out.

Figure 26–4 Multiple Address LAT Configuration: Two LANs with Version 5.3 LAT Nodes



26.3.2.2 Unsupported Configuration

When configuring a network to use an OpenVMS system running the LAT Version 5.3 software, *avoid* the configuration shown in Figure 26–5.

Figure 26–5 Unsupported Multiple Address LAT Configuration



Any configuration similar to this diagram will result in unpredictable results and may not function. In a network environment, LAT Version 5.1 and 5.2 nodes can have only a single logical LAN address. The configuration in Figure 26–5 violates this rule. The configuration shown in Figure 26–4 is a valid alternative.

26.3.2.3 Creating Logical LAT Links

The LAT software regards all paths as equal, point-to-point communications. The LAT software can support a maximum of eight LAN adapters simultaneously (and it is possible to connect all controllers to the same logical LAN).

To get the maximum coverage over possible path failures, each logical link should be created prior to setting the LAT node state to ON in SYS\$MANAGER:LAT\$SYSTARTUP.COM.

For example, if a system has one Ethernet adapter (device ESA0) with two FDDI adapters (FCA0 and FCB0) and the system manager chooses to run LAT over all adapters, the LAT\$SYSTARTUP.COM file would contain the following commands:

```

$!
$! Create each logical LAT link with a unique name and
$! unique LAN address (forced with /NODECNET).
$!
$ LCP CREATE LINK ETHERNET /DEVICE=ESA0 /NODECNET
$ LCP CREATE LINK FDDI_1 /DEVICE=FCA0 /NODECNET
$ LCP CREATE LINK FDDI_2 /DEVICE=FCB0 /NODECNET
$!
$! Turn on the LAT protocol.
$!
$ LCP SET NODE /STATE=ON

```

Caution

If the LATCP command SET NODE /STATE=ON is entered before the link is created, a random or default LAT\$LINK will be created on one of the LAN adapters. There is no way to predict which LAN adapter will be chosen (it is dependent on the system configuration). Therefore, all logical LAT links should be created before LAT is started.

Be sure each logical link is created with the /NODECNET qualifier. It will prevent the possibility of link creation failure if multiple adapters attempt to use the DECnet style address. Having more than one LAN adapter connected to the same logical LAN with the same address violates LAN conventions and will cause problems with LAT and other protocols.

Managing the LAT Software

26.3 Understanding LAT Configurations

It is possible to create logical LAT datalinks after the LAT protocol has been started. Any existing virtual circuit will attempt to find any new paths through the newly created logical datalink when it is ready for use. However, Compaq does not recommend that you create links at this point because during the time it takes existing virtual circuits to discover new paths through this newly created datalink, the virtual circuit may fail before the new paths are discovered.

26.3.2.4 Path Discovery

The OpenVMS LAT software uses a combination of the directory service and solicitation to obtain paths for each virtual circuit. To expedite path discovery at virtual circuit startup, Compaq recommends that you configure a system with multiple LAN adapters to maintain a LAT service and node database, by performing the following actions:

- Enabling outgoing LAT connections
- Using the same group code mask for User Groups and Service Groups

An OpenVMS system running with outgoing connections disabled and no service and node database is still capable of running with multiple paths for each virtual circuit. These paths must be discovered through the LAT solicitation process and will take longer (leaving the possibility for virtual circuit failure to occur before all paths have been discovered).

26.3.2.5 Modifying LAT Parameters

In the unlikely event of a path failure, it will take the OpenVMS LAT software time (which will vary depending on the number of adapters to which the remote node has access) to locate another working path. Therefore, Compaq strongly recommends that you modify the following LAT parameters on potential LAT master nodes:

- Retransmit limit - default value is 8. Set to the maximum number of LAN adapters times 8. For example, if an OpenVMS system on the LAN has 3 adapters, each LAT master node should have their retransmit limit set to 24 ($3 * 8$).
- Keepalive timer - default value is 20 seconds. While the default value may be sufficient in most circumstances, it may be necessary to increase this to 30 or 40 seconds.

Although it is possible to keep virtual circuits running through multiple adapters to LAT Version 5.1 or LAT Version 5.2 master nodes, there is still a possibility that the connections to these nodes may fail.

LAT Version 5.2 and LAT Version 5.1 master nodes do not have the ability to recognize multiple paths to LAT nodes that provide services. They can only communicate with such nodes through one remote address at a time. Therefore, if a LAN path failure occurs when a LAT master node running LAT Version 5.1 or Version 5.2 attempts to connect to a remote LAT Version 5.3 node providing services, the LAT Version 5.3 node might not discover this failure in time and the LAT master node may time out the connection. You can partially solve this problem by increasing the retransmit limit to as high a setting as possible.

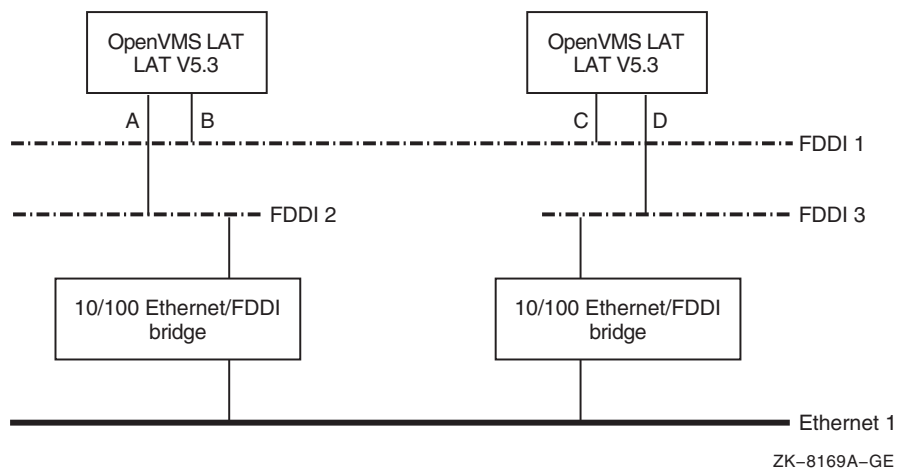
In addition, if a LAT Version 5.3 node providing services views the virtual circuit as completely idle during the primary path failure, no attempt will be made to use any of the alternate paths (because of the previously described LAT Version 5.2 and 5.1 limitation). Therefore, although multiple LAN adapters will work with older LAT implementations, you might still need to upgrade to Version

7.0 or higher of the OpenVMS operating system to obtain the LAT Version 5.3 protocol, which will correct this type of problem. (Note that this type of problem affects only those connections that are idle. An example of where this situation could arise is in an office environment if all users were to leave their systems at the same time, either at lunchtime or at the end of the workday.)

26.3.3 Large Buffers in Ethernet/FDDI Configurations

The OpenVMS LAT software will attempt to use *large* buffers over any virtual circuit that comes in over an FDDI controller. This feature can cause problems if an alternate virtual circuit path must go through an Ethernet. Figure 26–6 is an example of the configuration that can cause problems.

Figure 26–6 LAT FDDI Ring and Large Buffers



In this diagram, it is possible for the two OpenVMS systems to communicate using large packets through the path described by controllers B and C. *Large* packets may exceed 1500 bytes of data (the maximum Ethernet message can contain 1500 bytes of data). If the path described by controllers B and C were to fail, it will not be possible for communication to continue through the path described by A and D.

The path described by controllers A and D pass through an Ethernet LAN segment. The messages that are routed through the 10/100 bridges cannot be larger than the maximum Ethernet message. Problems can occur because the OpenVMS LAT software cannot always detect this kind of configuration.

There are two ways to prevent problems with the previously described configuration. The first and easiest option is to create a logical LAT link using an Ethernet adapter (if either system has an Ethernet LAN adapter). This will force the message size negotiation to be no larger than the maximum sized Ethernet message.

If neither system has an Ethernet controller (thus making the first option not possible), the second option is to override the use of large buffer support (which is enabled by default) by using the new LATCP command qualifier, `/[NO]LARGE_BUFFER`. For example:

```
$ MCR LATCP SET NODE/NOLARGE_BUFFER
```

Managing the LAT Software

26.3 Understanding LAT Configurations

Compaq recommends that you use the SET NODE/NOLARGE_BUFFER command after all logical LAT links have been created and before the LAT node has been turned on. For example, note the order of the commands in LAT\$SYSTARTUP.COM:

```
$!  
$! Create each logical LAT link with a unique name and  
$! unique LAN address (forced with /NODECNET).  
$!  
$ LCP CREATE LINK FDDI_1 /DEVICE=FCA0 /NODECNET  
$ LCP CREATE LINK FDDI_2 /DEVICE=FCB0 /NODECNET  
$!  
$! Don't use large buffer support (force packet  
$! sizes to be no larger than what Ethernet can  
$! support).  
$!  
$ LCP SET NODE /NOLARGE_BUFFER  
$!  
$! Turn on the LAT protocol.  
$!  
$ LCP SET NODE /STATE=ON
```

26.4 Understanding the LATCP Utility

The LAT Control Program (LATCP) utility configures and controls the LAT software on OpenVMS systems. LATCP commands let you stop and start the LAT driver (which implements the LAT protocol) and modify or display LAT characteristics of the OpenVMS node.

With LATCP, you can set up your system as a service node, which offers one or more resources (services) for access by users on other systems in the local area network (LAN).

You can also use LATCP to set up the system to allow its users to access services on other systems in the LAN. In this case, the system can act like a terminal server: it can manage multiple user sessions simultaneously for connections to services on other nodes.

You can use LATCP to set up your system to support incoming access only, outgoing access only, or both incoming and outgoing access. You can also set up your system so that it supports neither incoming nor outgoing access.

When you set up your system to support outgoing access, the LAT software manages a database of LAT services and nodes. The software builds the database when you enable outgoing access on your node. The software begins to collect **LAT service announcements**—multicast messages sent by LAT service nodes—and builds the database based on these service announcements. You can use LATCP to display the services and nodes in this database and to control the size of the database. Allow outgoing access on systems that can tolerate the additional overhead, such as standalone systems.

Use LATCP to perform the following actions:

- Specify operational characteristics for your node and its services
- Turn the state of the LAT port driver (LTDRIVER) on and off
- Display the status of LAT services and service nodes in the network
- Display the status of links created on your LAT node
- Display the status of your LAT node

- Show and zero LAT counters
- Create, delete, and manage LAT ports
- Recall previously entered LATCP commands so that you can execute them again without having to reenter them
- Create subprocesses so that you can execute DCL commands without exiting from LATCP

With the LAT protocol, you can set up LAT application ports on the local node so that users can access printers and other asynchronous devices that are connected to LAT terminal servers or service nodes on the LAN. The remote devices must be configured appropriately.

26.4.1 Invoking and Exiting LATCP

Enter the following command to invoke LATCP:

```
$ RUN SYSSYSTEM:LATCP
LATCP>
```

At the LATCP> prompt, you can enter LATCP commands. To exit LATCP, enter EXIT or press Ctrl/Z at the LATCP> prompt.

You can also execute a single LATCP command by using a DCL string assignment statement, as shown in the following example:

```
$ LCP := $LATCP
$ LCP SET NODE/STATE=ON
```

LATCP executes the SET NODE command and returns control to DCL.

26.4.2 LATCP Commands

Table 26–1 summarizes the LATCP commands.

Table 26–1 LATCP Commands

Command	Function
ATTACH	Transfers control from your current process to the specified process.
CREATE LINK	Creates LAT data links.
CREATE PORT	Creates a logical port on the local node.
CREATE SERVICE	Creates a service on a service node.
DEFINE/KEY	Assigns a command string to a function key on your keypad.
DELETE LINK	Deletes a LAT data link from a node.
DELETE PORT	Deletes an application port or dedicated port.
DELETE QUEUE_ENTRY	Deletes an incoming queued request from the local node.
DELETE SERVICE	Deletes a service on a service node.
EXIT	Returns the user to DCL command level.
HELP	Displays help text for LATCP commands.

(continued on next page)

Managing the LAT Software

26.4 Understanding the LATCP Utility

Table 26–1 (Cont.) LATCP Commands

Command	Function
RECALL	Recalls LATCP commands that you entered previously so that you can execute them again.
REFRESH	Refreshes your display screen, for example, after your display has been overwritten by output from some other source.
SCROLL	Allows you to retrieve information that has scrolled off the screen.
SET LINK	Modifies characteristics of LAT data links.
SET NODE	Specifies LAT characteristics for a node.
SET PORT	Maps a logical port on a node to either a remote device on a terminal server or a special application service on a remote LAT service node.
SET SERVICE	Changes service characteristics.
SHOW LINK	Displays the characteristics of links on your node.
SHOW NODE	Displays the characteristics of nodes.
SHOW PORT	Displays port characteristics.
SHOW QUEUE_ENTRY	Displays information about requests, or entries, queued on the local node.
SHOW SERVICE	Displays characteristics of LAT services known to your node.
SPAWN	Creates a subprocess.
ZERO COUNTERS	Resets the node counters, service counters, and link counters maintained by your node.

For detailed information about LATCP commands and qualifiers, refer to the *OpenVMS System Management Utilities Reference Manual*.

26.5 Starting Up the LAT Protocol

As system manager, you start up the LAT protocol and configure your node as a service node by executing the command procedure `SYS$STARTUP:LAT$STARTUP`. This procedure executes the following two procedures:

1. `LAT$CONFIG.COM`, to load the LAT terminal driver `LTDRIVER` and create the `LATACP` process
2. `LAT$SYSTARTUP.COM`, to execute LATCP commands that define LAT characteristics

How to Perform This Task

To make sure the LAT protocol is started each time the system boots, add a command to execute this procedure in the general-purpose, site-specific startup command procedure, described as follows. (See Section 5.2.1 for more detailed information about this command procedure, including the file specification used to identify it in your operating system.)

Managing the LAT Software

26.5 Starting Up the LAT Protocol

To set up your node as a LAT service node and start the LAT protocol software on your system each time the system boots, edit the general-purpose, site-specific startup command procedure to add the following line:

```
$ @SYS$STARTUP:LAT$STARTUP.COM
```

When the general-purpose, site-specific startup command procedure executes this command, it invokes LAT\$STARTUP.COM, which in turn invokes the LAT\$CONFIG and LAT\$SYSTARTUP command procedures.

You can append any of the following arguments to the command line that invokes LAT\$STARTUP to specify unique LAT characteristics for your node. The procedure will pass these arguments to LAT\$SYSTARTUP.COM to define the LAT characteristics you specify.

```
$ @SYS$STARTUP:LAT$STARTUP "P1" "P2" "P3" "P4" "P5"
```

Compaq recommends that you modify LAT\$SYSTARTUP.COM directly, rather than passing parameters in *P1* through *P5*. However, if you choose to use *P1* through *P5*, the arguments have the following meanings:

Argument	Format	Meaning
<i>P1</i>	Service-name	Name of the service. For clustered service nodes, use the cluster alias as the service name. For independent service nodes, use the DECnet node name. LAT\$SYSTARTUP.COM uses the argument <i>P1</i> to assign a service name to the node (with the LATCP CREATE SERVICE command).
<i>P2-P4</i>	Any of the following names:	LAT\$SYSTARTUP.COM uses the arguments to assign LAT node characteristics (with the LATCP SET NODE command).
	/IDENTIFICATION= "string"	Description of the node and its services that is advertised over the local area network (LAN). The default is the string defined by the logical name SYS\$ANNOUNCE. Make sure you include five sets of quotation marks around the identification string. For example: "/IDENTIFICATION=" - """"Official system center""""
	/GROUPS=(ENABLE= <i>group-list</i>)	Terminal server groups qualified to establish connections with the service node. By default, group 0 is enabled.
	/GROUPS=(DISABLE= <i>group-list</i>)	Removes previously enabled terminal server groups. If you are specifying the preceding qualifier to enable groups, you can combine the qualifiers into one, as shown in the example that follows this table.

Managing the LAT Software

26.5 Starting Up the LAT Protocol

Argument	Format	Meaning
P5	Any qualifiers valid with the CREATE SERVICE command	LAT\$SYSTARTUP.COM uses this argument to assign service characteristics with the LATCP CREATE SERVICE command. You can specify the /IDENTIFICATION, /LOG, and /STATIC_RATING qualifiers. Specify several qualifiers as shown in the following example: <pre>"/IDENTIFICATION=" - """"Official system node"""" - /STATIC_RATING=250"</pre>

Note that if you want to do any of the following LAT network tasks, you must edit LAT\$SYSTARTUP.COM (described in Section 26.6):

- Set up LAT printers.
- Create special application services.
- Set up the node to allow outgoing connections (to support the SET HOST/LAT command).

For a full description of LATCP commands and qualifiers, refer to the *OpenVMS System Management Utilities Reference Manual*.

Example

The following command creates the service OFFICE on the service node MOE, which is part of the OFFICE cluster (refer to Figure 26–1):

```
$ @SYS$STARTUP:LAT$STARTUP OFFICE
```

26.6 Customizing LAT Characteristics

To define special LAT characteristics for your node, edit the site-specific command procedure SYS\$MANAGER:LAT\$SYSTARTUP.COM. This command procedure contains LATCP commands that define LAT characteristics. LAT\$SYSTARTUP.COM is invoked when you execute the LAT\$STARTUP command procedure. As explained in Section 26.5, you typically execute LAT\$STARTUP.COM from the general-purpose, site-specific startup command procedure.

If you want your node to be a LAT service node that only supports incoming connections from interactive terminals, editing LAT\$SYSTARTUP.COM is not necessary. You can assign a service name and other characteristics by specifying parameters when you invoke the command procedure SYS\$STARTUP:LAT\$STARTUP, as described in Section 26.5.

However, you can edit LAT\$SYSTARTUP.COM to add LATCP commands that customize LAT characteristics for your node, for example:

Task	For More Information
Create more than one service	Section 26.6.1
Create logical ports for special application services and printers	Section 26.6.2
Enable queued incoming requests	Section 26.6.3

Managing the LAT Software

26.6 Customizing LAT Characteristics

Task	For More Information
Enable outgoing LAT connections to support the SET HOST/LAT command	Section 26.6.4
Tailor node characteristics ¹	Section 26.6.5

¹For example, to assign special service announcements or LAN links (using the SET NODE and SET LINK commands).

Caution

Do not edit the command procedures LAT\$STARTUP.COM and LAT\$CONFIG.COM. These are procedures supplied by Compaq to perform functions necessary for the LAT protocol to run correctly. Edit only LAT\$SYSTARTUP.COM to define LAT characteristics specific to your site.

If you edit LAT\$SYSTARTUP.COM, you should add only LATCP commands. In addition, you should conform to the order of commands in the template file SYS\$MANAGER:LAT\$SYSTARTUP.TEMPLATE. Section 26.6.5 provides a sample edited LAT\$SYSTARTUP procedure. The *OpenVMS System Management Utilities Reference Manual* contains full descriptions of all the LATCP commands you can include in LAT\$SYSTARTUP.COM.

26.6.1 Creating Additional Services

The LAT\$SYSTARTUP.COM procedure provided by Compaq creates one service. This can be a primary service, one through which users can access the general computing environment. It can also be a special application service, such as a data entry program or an online news service.

You can also create a limited service with a fixed number of LTA devices, as described in Section 26.6.2.3.

The LAT\$SYSTARTUP.COM procedure creates the service with the same name as that of your node, unless you specify a unique service name as an argument to the @SYS\$STARTUP:LAT\$STARTUP.COM command, as explained in Section 26.5.

How to Perform This Task

To create services in addition to the one provided in LAT\$SYSTARTUP.COM, use the CREATE SERVICE commands, which you can add to LAT\$SYSTARTUP.COM. Note that if you create an application service, Compaq recommends that you assign the name of the application program. For more information about the LATCP command CREATE SERVICE, refer to the *OpenVMS System Management Utilities Reference Manual*.

Example

The following example creates the application service NEWS on the local node:

```
$ LCP := $LATCP
$ LCP CREATE SERVICE /APPLICATION NEWS
```

Managing the LAT Software

26.6 Customizing LAT Characteristics

26.6.2 Setting Up Ports

The LAT\$SYSTARTUP.COM file provided by Compaq includes sample commands to create logical ports on the service node and associates them with physical ports or services on the terminal server node. These ports can be used for application services and remote printers.

How to Perform This Task

To create ports, enable the sample commands by removing the exclamation points (!) that precede them in the LAT\$SYSTARTUP.COM file, or add similar CREATE PORT and SET PORT commands to that file to meet your needs. For information about the LATCP commands CREATE PORT and SET PORT, refer to the *OpenVMS System Management Utilities Reference Manual*.

Note

Compaq strongly recommends that you create application ports and dedicated ports *after* the LATCP command SET NODE/STATE=ON is executed. This minimizes nonpaged pool memory usage and eliminates the possibility of creating duplicate ports.

Note that you may encounter the following error when attempting to create an application port (with a command such as LCP CREATE PORT LTA5001:/APPLICATION, for example):

```
%LAT-W-CMDERROR, error reported by command executor
-SYSTEM-F-DUPLNAM, duplicate name
```

This error indicates that the LAT application port you are trying to create is already created by some other application. This application could be LATCP itself (LATCP's port, LATCP\$MGMT_PORT, is used to communicate with LTDRIVER).

To avoid this error, make sure the SET NODE/STATE=ON command is executed before any commands that create application ports or dedicated ports. You can also use the LATCP command SET NODE/DEVICE_SEED. For more information about the SET NODE/DEVICE_SEED command, refer to the *OpenVMS System Management Utilities Reference Manual*.

26.6.2.1 Setting Up Printers

If you set up a port for a printer, you must also perform the following tasks:

1. Create a spooled output queue for the printer.
2. Add a command to start the queue to the startup command procedure that starts your queues, or to the general-purpose, site-specific startup command procedure.

These tasks are described in Chapter 14.

26.6.2.2 Setting Up Special Application Services

To establish a special application service, include the /DEDICATED qualifier when defining a LAT port. The application program to which the service connects must define the same dedicated port. For example, the following commands set up ports for an application service called NEWS:

```
$ LCP ::= $LATCP
$ LCP CREATE PORT LTA333:/DEDICATED
$ LCP SET PORT LTA333:/SERVICE=NEWS
```

Managing the LAT Software

26.6 Customizing LAT Characteristics

Before application services can be available to user terminals on the LAT network, you must start the application program. You usually add commands to SYLOGIN.COM to do this.

26.6.2.3 Setting Up Limited Services

Application services with dedicated ports allow you to create a predetermined number of LTA devices (LAT terminals, for example) that are under the control of a process supplied by the system. In that environment, however, the user cannot log in to the service because no way exists for dedicated LTA devices to run the system login image (LOGINOUT.EXE).

You can create a *limited* service that allows users to log in to a predetermined number of LTA devices associated with that limited service. When all those devices are in use, the LAT software will reject additional connection requests to that service, as indicated by “service in use” error messages. Creating a limited service in this way allows you to control the number of LAT users on your system. (Note, however, that you cannot control which LTA device will be assigned when a user connects to the limited service.)

The following example sets up a limited service with two predetermined LTA devices:

```
$ LCP := $LATCP
$ LCP CREATE SERVICE /LIMITED RESTRICTED
$ LCP CREATE PORT LTA100 /LIMITED
$ LCP CREATE PORT LTA101 /LIMITED
$ LCP SET PORT LTA100 /SERVICE=RESTRICTED
$ LCP SET PORT LTA101 /SERVICE=RESTRICTED
```

When a user attempts to connect to the limited service named RESTRICTED, the LAT software will choose either LTA100 or LTA101 (whichever is available first) and complete the user connection. The user can then log in to that system. If another user connects to the service, that second connection attempt will be assigned to the remaining LTA device. The user can then log in to that second system. When the two devices associated with the limited service named RESTRICTED are both in use, any subsequent attempts to connect to that service will be rejected, as indicated by the “service in use” error message.

When a user logs out of the system (LTA100 or LTA101), that LTA device is *not* deleted. Instead, it is reset to accept the next connection request to the limited service.

26.6.3 Queuing Incoming Requests

By default, incoming requests to limited or application services are queued. This means that if you attempt to connect to a limited or application service (by using a terminal server port with forward queuing enabled or by entering the DCL command SET HOST/LAT/QUEUE), the LAT software will queue, rather than reject, this connection request if the service has no available ports.

How to Perform This Task

You can set up and manage a service that queues incoming connect requests as follows:

- Use the LATCP command SHOW SERVICE to determine whether the service has queuing enabled or disabled.
- If queuing is disabled, use the SET SERVICE /QUEUED command to enable queuing.

Managing the LAT Software

26.6 Customizing LAT Characteristics

- Use the SET NODE /QUEUE_LIMIT=*n* command on the local node to control the number of free queue slots (where *n* is between 0 and 200).
- Use the SHOW QUEUE_ENTRY [queue-entry-id] command to view entries in the local queue.
- Use the DELETE QUEUE_ENTRY [queue-entry-id] command to delete an entry from the local queue.

Refer to the *OpenVMS System Management Utilities Reference Manual* for more detailed descriptions of the LATCP commands and qualifiers you use to support queued requests.

Example

The following example shows how to enable queuing on your system:

```
$ LCP :== $LATCP
$ LCP SET SERVICE /QUEUE
```

Note

If a system is configured to handle queued connect requests, that system *must* be set up as follows to avoid possible queue connection failures:

- Incoming and outgoing connections must be enabled.
 - User group codes and service group codes must be identical.
-

26.6.4 Enabling Outgoing LAT Connections

By default, outgoing LAT connections are disabled on a node. To allow users to use the SET HOST/LAT connection to establish LAT connections from the node, you must edit LAT\$SYSTARTUP.COM to enable outgoing connections. For more details on using the SET HOST/LAT command for outgoing LAT connections, refer to the description of that command in the *OpenVMS DCL Dictionary*.

Commands to enable outgoing connections are included in the LAT\$SYSTARTUP.COM file provided by Compaq. Enable the command of your choice by removing the exclamation point (!) that precedes it, or add a similar command to meet your needs. For more information, refer to the /CONNECTIONS and /USER_GROUPS qualifiers to the LATCP command SET NODE in the *OpenVMS System Management Utilities Reference Manual*.

To attain optimal SET HOST/LAT performance and forward port performance, you should set the system parameter TTY_ALTYPAMD to 1500 and reboot.

If you want to set up your node only as a service node with incoming connections enabled, editing LAT\$SYSTARTUP.COM is not necessary. However, you might edit LAT\$SYSTARTUP.COM to do one or more of the following tasks:

- Create more than one service on a node
- Create special application services
- Set up LAT printers
- Enable outgoing LAT connections (to allow a node to act as a terminal server node)
- Tailor node characteristics; for example, to assign special service announcements or connections to the LAN

26.6.5 Sample Edited LAT\$SYSTARTUP.COM Procedure

The following example is a sample of an edited LAT\$SYSTARTUP.COM procedure that creates services, creates and sets ports, and sets nodes to allow incoming and outgoing connections.

```

$!
$! LAT$SYSTARTUP.COM -- LAT Startup Commands Specific to Site
$!
$! Use this command procedure to customize the LAT characteristics for
$! the local node. These commands, which should serve as examples,
$! will set up a LAT service name SYS$NODE and default identification
$! SYS$ANNOUNCE. The LAT service name and identification will default
$! to SYS$NODE and SYS$ANNOUNCE unless you specify a service name and
$! identification as arguments to the command line that invokes
$! LAT$STARTUP.COM:
$!          $ @SYS$STARTUP:LAT$STARTUP
$!
$! You can specify other node and service characteristics (such as group
$! codes) as arguments to this command line, as shown below.
$!
$!      Argument      Function
$!      -----      -
$!
$!          P1          Name of the service to be created. If not supplied, a
$!                      service will be created with the same name as the node.
$!
$!          P2,P3,P4    Parameters and qualifiers to the SET NODE command.
$!
$!          P5          Parameters and qualifiers to the SET SERVICE command.
$!                      P5 is only used if P1 is specified. More than one
$!                      argument may be supplied by enclosing the string in
$!                      quotes.
$!
$!      Example: $ @SYS$STARTUP:LAT$STARTUP HAWK "/IDENTIFICATION=" -
$!              "Development node"
$!
$!      Please review and edit this file for possible additions and deletions
$!      that you wish to make. Future software updates will not overwrite the
$!      changes made to this file.
$!
$ required_privileges = "OPER"
$ prev_privs = f$setprv(required_privileges)
$ if .not. f$privilege(required_privileges) then goto no_privileges
$ lcp := $latcp
$!
$! ----- Modify Node Characteristics -----
$!
$ lcp set node 'p2' 'p3' 'p4'
$!
$! Some examples:
$!
$! ** Allow incoming connections only
$!
$! lcp set node /connections=incoming /groups=(enable=(12,40,43,73),disable=0)
$! lcp set node /connections=incoming /groups=enable=(0-255)
$!
$! LCP SET NODE /CONNECTIONS=INCOMING /GROUPS=(ENABLE=(12,40,43,73),DISABLE=0)
$!
$! ** Allow outgoing connections only
$!
$! lcp set node /connections=outgoing /user_groups=enable=(24,121-127)
$! lcp set node /connections=outgoing /user_groups=(enable=0-255) /node_limit=50
$!
$! ** Enable incoming and outgoing connections
$!

```

Managing the LAT Software

26.6 Customizing LAT Characteristics

```
#! lcp set node /connections=both /group=enable=(43,73) /user=enable=(44,56)
#! lcp set node /connections=both /group=enable=(0-255) /user=enable=(0-255)
#!
#!
#! ----- Modify Service Characteristics -----
#!
$ if pl .eqs. ""
$ then
$   lcp create service
$ else
$   lcp create service 'p1' 'p5'
$ endif
#! ----- Start LAT Protocol -----
#!
$ lcp set node /state=on
#!
#!
#! ----- Create and Map Ports -----
#!
$! Some examples:
$!
$! lcp create port lta101: /dedicated
$! lcp create port lta102: /application
$! lcp create port lta103: /application
$! lcp create port /nolog/logical=(name=ln03$mgmt, table=system, mode=executive)
$
$ LCP CREATE PORT LTA1: /NOLOG
$ LCP CREATE PORT LTA20: /NOLOG
$
$! lcp set port lta101: /dedicated /service=graphics
$! lcp set port lta102: /node=server_1 /port=port_1
$! lcp set port lta103: /node=server_2 /service=laser
$! lcp set port ln03$mgmt: /node=server_3 /service=ln03_printers
$!
$ LCP SET PORT LTA1: /APPLICATION/NODE=TERM_SERVER_1 /PORT=PORT_6
$ LCP SET PORT LTA20: /APPLICATION/NODE=TERM_SERVER_2 /PORT=PORT_6
$!
$exit:
$ prev_privs = f$setprv(prev_privs)
$ exit
$!
$no_privileges:
$ write sys$output "Insufficient privileges to execute LATCP commands."
$ write sys$output "Requires ",required_privileges," privileges."
$ goto exit
```

26.7 Managing the LATACP Database Size

On OpenVMS nodes, another component of the LAT software, the LAT Ancillary Control Process (LATACP), maintains the database of available nodes and services. The nodes and services can be those that are multicast from remote LAT nodes, or they can consist of the local node and one or more local services that you create on your own system. The maximum size of this database is dependent on the value of the system parameter CTLPAGES.

After you enter a LATCP command, you might get the following response:

```
%LAT-W-CMDERROR, error reported by command executor
-LAT-F-ACPNOCTL, insufficient resources - ACP CTL/P1 space limit reached
```

Managing the LAT Software

26.7 Managing the LATACP Database Size

If so, this signifies that the database size has reached the CTLPAGES limit. You can correct the situation in one of three ways:

- Reduce the size of the database by reducing the node limit. Use the LATCP command SHOW NODE to display the node limit; use the command SET NODE/NODE_LIMIT to change it. For more information, refer to the *OpenVMS System Management Utilities Reference Manual*.
- Reduce the size of the database by reducing the user group codes that are enabled on the node. Use the LATCP command SHOW NODE to display the enabled user group codes; use the command SET NODE/USER_GROUPS=DISABLE to disable some of them. For more information, refer to the *OpenVMS System Management Utilities Reference Manual*.

If you choose this step, you must also edit your startup procedures to change the user groups that are enabled each time the system reboots. For more information, see Section 26.6.

- Extend the size of the database by increasing the value of the system parameter CTLPAGES. As a general rule, note that every unit of CTLPAGES that you increase is roughly equivalent to six additional nodes or services that will be stored in the database.

After you change CTLPAGES, you must reboot the system for the changed value to take effect. Make sure you add the increased value of CTLPAGES to the AUTOGEN parameter file MODPARAMS.DAT. For more information about changing values of system parameters, see Section 15.2.

Managing DECdtm Services

This chapter describes what you must do if you want to run software that uses DECdtm services, such as ACMS, DECintact, Oracle Rdb, and RMS Journaling.

Note

On OpenVMS Alpha systems, unpredictable results can occur if DECdtm services are used in a multithreaded environment. Do not make calls to DECdtm services in kernel threads other than the initial thread because much of the work performed by DECdtm uses the context of the calling process.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Planning transaction logs	Section 27.2
Planning for a DECnet-Plus network	Section 27.3
Creating transaction logs	Section 27.4
Monitoring transaction performance	Section 27.5
Checking whether a transaction log is too small	Section 27.6
Changing the size of a transaction log	Section 27.7
Moving a transaction log	Section 27.8
Dismounting a disk	Section 27.9
Adding a node	Section 27.10
Removing a node	Section 27.11
Disabling DECdtm services	Section 27.12
Enabling DECdtm services	Section 27.13
Using the XA Gateway (Alpha only)	Section 27.14

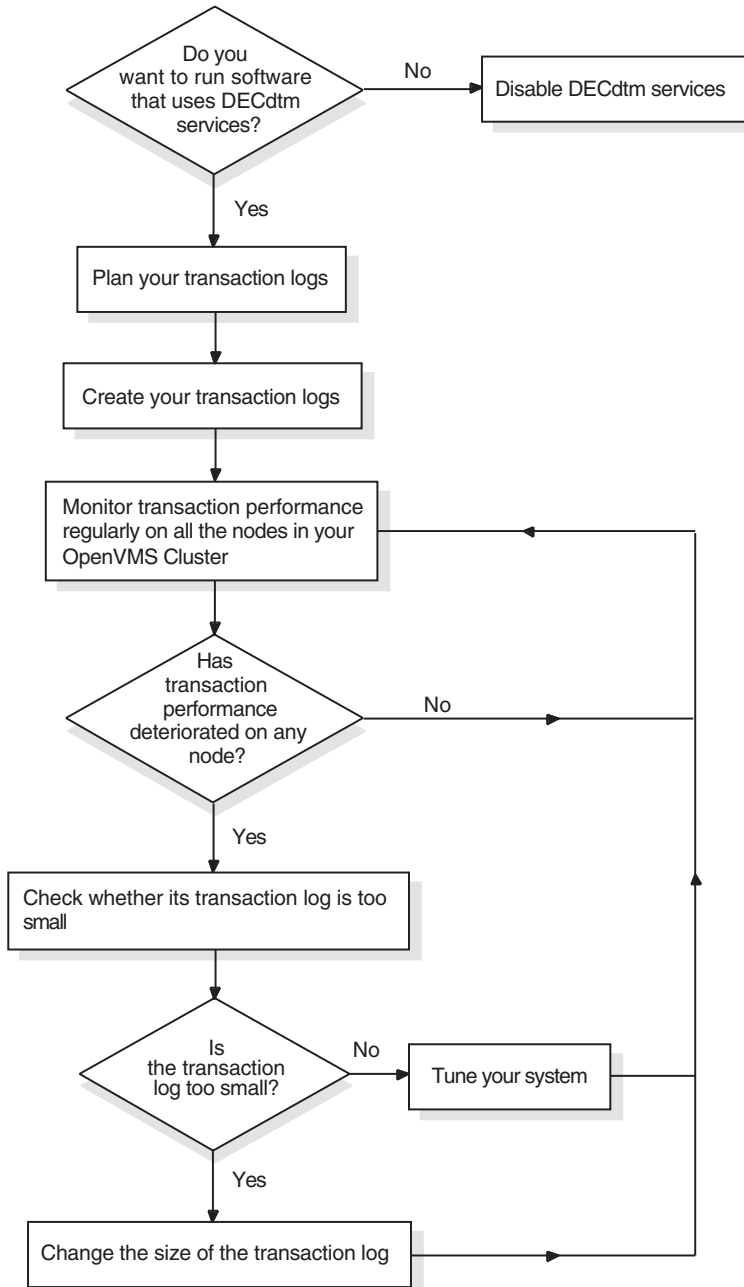
The map in Figure 27–1 shows the tasks, and the order in which to do them.

This chapter explains the following concepts:

Concept	Section
Understanding transaction logs	Section 27.1
Understanding transaction groups	Section 27.3.2.2

Managing DECdtm Services

Figure 27-1 Managing DECdtm Services



ZK-5154A-GE

27.1 Understanding Transaction Logs

A **transaction log** is a file that stores information about DECdtm transactions performed on a node. It is of file type .LM\$JOURNAL.

Before a node can execute DECdtm transactions, you must create a transaction log for the node. In an OpenVMS Cluster, create a transaction log for each node in the cluster. Use the Log Manager Control Program (LMCP) utility to create and manage transaction logs.

DECdtm services use the logical name SYS\$JOURNAL to find transaction logs. You must define SYS\$JOURNAL to point to the directories that contain transaction logs.

27.2 Planning Transaction Logs

The size and location of a transaction log can affect transaction performance. Before you create a transaction log, decide the size and location of the transaction log.

Later, you can change the size of a transaction log, or move it. However, careful planning at this stage reduces the need for future changes.

This section describes:

Task	Section
Deciding the size of a transaction log	Section 27.2.1
Deciding the location of a transaction log	Section 27.2.2

27.2.1 Deciding the Size of a Transaction Log

When you create a transaction log, you can specify its size. The default size is 4000 blocks; this gives acceptable performance on most systems.

If you know the expected rate of transactions, Compaq suggests the following formula to calculate the transaction log size:

$$size = 40 * rate$$

where:

size is the size of the transaction log in blocks.

rate is the average number of transactions executed per second.

If you do not know the rate of transactions, accept the default size of 4000 blocks.

Managing DECdtm Services

27.2 Planning Transaction Logs

27.2.2 Deciding the Location of a Transaction Log

If possible, choose a disk that is:

Fast	Achieve speed by using a high-performance disk, such as a solid-state disk, that is not heavily used.
Highly available	Achieve high availability by having multiple access paths to the data. In an OpenVMS Cluster, use a disk that can be accessed by the other nodes in the cluster. This ensures that if one node fails, transactions running on other nodes are not blocked.
Reliable	Achieve reliability by keeping multiple copies of the data. Using a shadowed disk is more reliable than using a nonshadowed disk, but may be slower because transaction logs are almost exclusively write-only.

You may need to choose between speed and either availability or reliability. For example, if the node is a workstation, you may choose to sacrifice speed for availability and reliability by putting the node's transaction log on a shadowed HSC-based disk, instead of on a faster disk attached to the workstation.

In a cluster environment, try to distribute the transaction logs across different disks. Having more than one transaction log on a disk can lead to poor transaction performance.

Note

Make sure that the disk has enough contiguous space to hold the transaction log. A discontinuous transaction log leads to poor transaction performance.

27.3 Planning for a DECnet-Plus Network

This section contains the following information to help you plan for using DECdtm in a DECnet-Plus network:

- Planning your DECnet-Plus namespace
- Planning SCSNODE names in your DECnet-Plus network

27.3.1 Planning Your DECnet-Plus Namespace

DECdtm does not support multiple DECnet-Plus namespaces.

This means that if you want to use software that uses DECdtm services, you cannot use both a local namespace and a DECdns namespace.

27.3.2 Planning SCSNODE Names in Your DECnet-Plus Network

SCSNODE is a system parameter that defines the name of the computer. You must follow certain rules when choosing SCSNODE names if you have a DECnet-Plus network and you want to perform DECdtm transactions that span either different OpenVMS Clusters or different standalone computers.

27.3.2.1 Rules for SCSNODE Names

If you have a DECnet-Plus network and want to perform DECdtm transactions that span different OpenVMS Clusters or different standalone computers, you must make sure that your SCSNODE names obey the following rules:

- The SCSNODE name of each computer in a transaction group must be different from:
 - The SCSNODE names of the other computers in the transaction group; SCSNODE names must be unique within a transaction group
 - The DECnet simple names of other computers on the same local root
 - The DECnet synonyms of the other computers in the entire network
- For an explanation of transaction groups, see Section 27.3.2.2.
- If a computer is in an OpenVMS Cluster, its SCSNODE name must also be different from:
 - The DECnet simple names of the other computers in the same cluster
 - The DECnet simple names of computers on the same local root as other cluster members

27.3.2.2 Understanding Transaction Groups

A **transaction group** is a group of computers involved in DECdtm transactions whose SCSNODE names must obey the rules described in Section 27.3.2.1.

A transaction group conforms to the following guidelines:


- Each computer belongs to no more than one transaction group.
- All the computers in an OpenVMS Cluster belong to the same transaction group.
- If a single transaction spans computers A and B, then computers A and B belong to the same transaction group.

Figure 27–2 shows an example of a transaction group.

Managing DECdtm Services

27.3 Planning for a DECnet-Plus Network

Figure 27–2 Transaction Group



ZK-6302A-GE

All nine computers shown in the figure are in the same transaction group because:

- A transaction spans a computer in cluster FRED and a computer in cluster BILL. This means that the four computers in cluster FRED and the four computers in cluster BILL are in the same transaction group.
- A transaction spans standalone computer TOM and a computer in cluster BILL. This means that the standalone computer TOM is in the same transaction group as the computers in cluster BILL.

27.4 Creating Transaction Logs

Before a node can perform DECdtm transactions, you must create a transaction log for the node. In an OpenVMS Cluster environment, create a transaction log for each node.

Caution

Removing a node from a cluster after you have created the transaction logs can lead to data corruption. For instructions on how to remove a node safely, see Section 27.11.

How to Perform This Task

1. For each node, decide the size and location of the transaction log, using the guidelines in Section 27.2. Remember that the disks must have enough contiguous space to hold the transaction logs.
2. If you are in a cluster environment, make sure that the disks on which you want to create the transaction logs are mounted clusterwide.
3. Decide in which directories you want to create the transaction logs. You may want to create new directories for the transaction logs.
4. Define SYS\$JOURNAL to point to the directories in which you want to create the transaction logs:

```
DEFINE/SYSTEM/EXECUTIVE_MODE SYS$JOURNAL dirspec[,...]
```

where *dirspec* is the full specification of a directory in which you want to create one or more transaction logs. List all the directories that will contain transaction logs. You can list the directories in any order.

In a cluster environment, use SYSMAN to define SYS\$JOURNAL clusterwide.

5. Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to include the SYS\$JOURNAL definition.

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

6. Create one transaction log for each node, using LMCP's CREATE LOG command:

```
CREATE LOG [/SIZE=size] dirspecSYSTEM$node.LM$JOURNAL
```

where:

size is the size of the transaction log in blocks. By default, the size of the transaction log is 4000 blocks.

dirspec is the full specification of the directory in which you want to create the transaction log.

node is the name of the node.

7. Make sure DECdtm services are enabled as follows:

Step	Action
a.	Check whether the logical SYS\$DECDTM_INHIBIT is defined: <pre>\$ SHOW LOGICAL SYS\$DECDTM_INHIBIT</pre>
b.	Is SYS\$DECDTM_INHIBIT defined? Yes DECdtm services are disabled. Enable DECdtm services by following the instructions in Section 27.13. No DECdtm services are enabled.

Example

This example shows how to create transaction logs in an OpenVMS Cluster that consists of two nodes whose SCSNODE names are BLUE and RED. Neither node has a node-specific version of SYLOGICALS.COM.

Managing DECdtm Services

27.4 Creating Transaction Logs

Decide the size and location of the transaction logs:

Node	Size of Log (in Blocks)	Disk
BLUE	5000	DUA1
RED	4000	DUA2

Mount the disks clusterwide:

```
$ MOUNT/CLUSTER/SYSTEM DUA1: LOG1
$ MOUNT/CLUSTER/SYSTEM DUA2: LOG2
```

Create directories for the transaction logs:

```
$ CREATE/DIRECTORY DISK$LOG1:[LOGFILES]
$ CREATE/DIRECTORY DISK$LOG2:[LOGFILES]
```

Define SYS\$JOURNAL:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEFINE/SYSTEM/EXECUTIVE MODE SYS$JOURNAL -
SYSMAN> DISK$LOG1:[LOGFILES], DISK$LOG2:[LOGFILES]
SYSMAN> EXIT
```

Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to include the following line:

```
$ !
$ DEFINE/SYSTEM/EXECUTIVE_MODE SYS$JOURNAL DISK$LOG1:[LOGFILES], -
DISK$LOG2:[LOGFILES]
$ !
```

Create the transaction logs:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CREATE LOG/SIZE=5000 DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL
LMCP> CREATE LOG DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$JOURNAL
LMCP> EXIT
```

Make sure DECdtm services are enabled:

```
$ SHOW LOGICAL SYS$DECDTM_INHIBIT
%SHOW-S-NOTRAN, no translation for logical name SYS$DECDTM_INHIBIT
```

SYS\$DECDTM_INHIBIT is undefined, so DECdtm services are enabled.

27.5 Monitoring Transaction Performance

Changes to your system, such as increase in work load, can affect transaction performance. Once a month, monitor transactions on the node to make sure that transaction performance has not deteriorated. In an OpenVMS Cluster environment, monitor transaction performance on all the nodes in the cluster.

How to Perform This Task

1. Monitor transactions, using the MONITOR TRANSACTION command:

```
MONITOR TRANSACTION/SUMMARY[=file-spec]/ENDING=time/NODE=(nodename,...)
```

where:

Managing DECdtm Services

27.5 Monitoring Transaction Performance

<i>file-spec</i>	is the file specification of the summary file. Information about transactions is summarized and recorded in the summary file. If you omit the file specification, the information is recorded in MONITOR.SUM in your default directory.
<i>time</i>	is the time that the monitoring session ends.
<i>nodename</i>	is the name of a node. In an OpenVMS Cluster, list all the nodes in the cluster.

For the best results, monitor transactions for a day at a time.

You can monitor transactions in batch mode by including the MONITOR TRANSACTION command in a command procedure.

For a full description of the MONITOR TRANSACTION command, refer to the *OpenVMS System Management Utilities Reference Manual*.

2. Examine the summary file.

The summary file contains values for a number of different data items. Note the following values for each node:

- Average end rate. This is the average number of transactions completed per second.
- Average completion rates. These are the average numbers of transactions completed in the following times:

- Less than 1 second
- Between 1 and 2 seconds
- Between 2 and 3 seconds
- Between 3 and 4 seconds
- Between 4 and 5 seconds
- More than 5 seconds

Keep a note of these values.

3. Compare the results from this monitoring session with the results from previous sessions.

For the same work load, the rate and duration of transactions should remain about the same. Indications of performance deterioration are:

- The average end rate decreases
- The average duration increases

To find out whether the average duration of transactions has increased, compare the average completion rates. If a greater proportion of the transactions takes longer to complete, the average duration of transactions has increased.

Note any trends over a number of monitoring sessions. Variations from one monitoring session to the next are probably due to variations in work load.

If you suspect that transaction performance has deteriorated on any node, check whether its transaction log is too small (see Section 27.6).

If the transaction log is big enough, but transaction performance still deteriorates, tuning the system might be necessary. Refer to the *OpenVMS Performance Management* for information about tuning your system.

Managing DECdtm Services

27.5 Monitoring Transaction Performance

Example

This example shows how to monitor transaction performance on an OpenVMS Cluster that consists of two nodes whose SCSNODE names are BLUE and RED.

Monitor transactions on nodes BLUE and RED for one day:

```
$ MONITOR TRANSACTION/SUMMARY=DISK$LOG1:[LOGFILES]TRANSACTIONS.SUM -
_ $ /ENDING="+1-"/NODE=(BLUE,RED)
```

Examine the summary file:

```

DISTRIBUTED TRANSACTION STATISTICS
      on node BLUE          From: 16-OCT-2000 14:23:51
      SUMMARY              To:  17-OCT-2000 14:23:51
      CUR          AVE      MIN      MAX
Start Rate          49.02    43.21    31.30    49.02
Prepare Rate        48.70    43.23    30.67    48.70
One Phase Commit Rate 0.00    0.00    0.00    0.00
Total Commit Rate   48.70    43.19    31.30    48.70
Abort Rate          0.00    0.00    0.00    0.00
End Rate            48.70    43.19    31.30    48.70
Remote Start Rate   0.00    0.00    0.00    0.00
Remote Add Rate     0.00    0.00    0.00    0.00

Completion Rate  0-1      21.42    13.57    0.63    21.42
  by Duration   1-2      25.97    29.15    24.59    33.87
  in Seconds    2-3       1.29    0.47    0.00    4.47
                3-4       0.00    0.00    0.00    0.00
                4-5       0.00    0.00    0.00    0.00
                5+       0.00    0.00    0.00    0.00

```

SUMMARIZING

```

DISTRIBUTED TRANSACTION STATISTICS
      on node RED          From: 16-OCT-2000 14:23:52
      SUMMARY              To:  17-OCT-2000 14:23:52

```

.
.
.

Make a note of the following values:

- Average end rate.
For node BLUE, the average end rate is 43.19 transactions per second.
- Average completion rates.
For node BLUE, the average completion rates are as follows:
13.57 transactions completed in 0 to 1 seconds
29.15 transactions completed in 1 to 2 seconds
0.47 transactions completed in 2 to 3 seconds

Compare the results from this monitoring session to those of previous sessions:

Session	End Rate	Completion Rates		
		0-1 Secs	1-2 Secs	2-3 Secs
June	42.13	12.98	28.13	1.02
July	38.16	10.35	25.80	2.01
August	43.19	13.57	29.15	0.47

The results for node BLUE show no signs of deteriorating performance.

27.6 Checking Whether a Transaction Log Is Too Small

If transaction performance has deteriorated on a node, check whether its transaction log is too small.

Section 27.5 describes how to find out whether transaction performance has deteriorated.

How to Perform This Task

1. Log in to the node that the transaction log belongs to.
2. Check how many times the transaction log has stalled, using LMCP's SHOW LOG/CURRENT command:

```
$ RUN SYSSYSTEM:LMCP
LMCP> SHOW LOG/CURRENT
```

Note the number of checkpoints and stalls displayed by this command.

3. Wait for five minutes, then repeat the SHOW LOG/CURRENT command. Note the number of checkpoints and stalls again.
4. Compare the information from the SHOW LOG/CURRENT commands:
If the number of checkpoints has not changed, wait until the system is busier, then try this task again.

If the number of checkpoints has increased, and the number of stalls has increased by more than one, the transaction log is too small.

5. If the transaction log is too small, increase its size. For information about how to change the size of a transaction log, see Section 27.7.

Example

This example shows how to check whether node BLUE's transaction log is too small.

Log in to node BLUE. Then check how many times the transaction log has stalled:

```
$ RUN SYSSYSTEM:LMCP
LMCP> SHOW LOG/CURRENT
```

```
Checkpoint starts/ends      2464/2464
Stall starts/ends          21/21
Log status: no checkpoint in progress, no stall in progress.
```

The number of checkpoints is 2464, and the transaction log has stalled 21 times.

Wait for five minutes, then repeat the SHOW LOG/CURRENT command:

```
LMCP> SHOW LOG/CURRENT
```

```
Checkpoint starts/ends      2514/2514
Stall starts/ends          28/28
Log status: no checkpoint in progress, no stall in progress.
```

The number of checkpoints has increased since the previous reading, and the transaction log has now stalled 28 times, an increase of 7. This means that the transaction log is too small.

Managing DECdtm Services

27.7 Changing the Size of a Transaction Log

27.7 Changing the Size of a Transaction Log

To determine if changing the size of a transaction log is necessary, see Section 27.6.

How to Perform This Task

Caution

Follow all the steps carefully. Taking shortcuts can lead to data corruption.

1. Log in to the node that the transaction log belongs to.
2. Find out which directory the transaction log is in, using LMCP's SHOW LOG command:

```
SHOW LOG SYSTEM$node.LM$JOURNAL
```

where *node* is the name of the node that the transaction log belongs to.

3. Rename the transaction log:

```
RENAME dirspecSYSTEM$node.LM$JOURNAL dirspecSYSTEM$node.LM$OLD
```

where:

dirspec is the full specification of the directory containing the transaction log.

node is the name of the node that the transaction log belongs to.

4. Can you stop all the software that uses DECdtm services without shutting down any nodes?

Yes Close the transaction log as follows:

Step	Action
------	--------

a. Stop all the software that uses DECdtm services.

b. Close the transaction log using LMCP's CLOSE LOG command:

```
$ RUN SYSSYSTEM:LMCP
LMCP> CLOSE LOG
```

The CLOSE LOG command closes the transaction log and stops the DECdtm TP_SERVER process. The command fails if any software is using DECdtm services.

c. Did the CLOSE LOG command succeed?

Yes Restart the TP_SERVER process:

```
$ @SYS$STARTUP:DECDTM$STARTUP.COM
```

No Wait for 30 seconds, then repeat steps 4b and 4c.

No Close the transaction log by rebooting the node. Log in to the node when it has rebooted.

5. Change the size of the transaction log, using LMCP's CONVERT LOG command:

```
CONVERT LOG/SIZE=size dirspecSYSTEM$node.LM$OLD dirspecSYSTEM$node.LM$JOURNAL
```


Managing DECdtm Services

27.7 Changing the Size of a Transaction Log

where:

size is the new size of the transaction log in blocks.
dirspec is the full specification of the directory containing the transaction log.
node is the name of the node that the transaction log belongs to.

6. If you stopped the software that uses DECdtm services in step 4, restart the software.

7. Delete the old transaction log:

```
DELETE dirspecSYSTEM$node.LM$OLD;
```

where:

dirspec is the full specification of the directory containing the old transaction log.
node is the name of the node that the transaction log belongs to.

Example

This example shows how to change the size of node RED's transaction log to 6000 blocks. Node RED is in an OpenVMS Cluster, and its transaction log is in DISK\$LOG2:[LOGFILES].

Log in to node RED. Find out which directory RED's transaction log is in, then rename the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> SHOW LOG SYSTEM$RED.LM$JOURNAL
Directory of DISK$LOG2:[LOGFILES]
SYSTEM$RED.LM$JOURNAL;1
Total of 1 file.
LMCP> EXIT

$ RENAME DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$JOURNAL -
_ $ DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$OLD
```

Stop all software that uses DECdtm services. Then close the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CLOSE LOG
Transaction log closed, TP_SERVER process stopped
LMCP> EXIT
```

Restart the TP_SERVER process:

```
$ @ SYS$STARTUP:DECDTM$STARTUP.COM
```

Change the size of the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CONVERT LOG/SIZE=6000 DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$OLD -
_ LMCP> DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$JOURNAL
Log file DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$JOURNAL;1 created.
Log file DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$OLD converted.
LMCP> EXIT
```

Restart the software that uses DECdtm services.

Delete the old transaction log:

```
$ DELETE DISK$LOG2:[LOGFILES]SYSTEM$RED.LM$OLD;
```

Managing DECdtm Services

27.8 Moving a Transaction Log

27.8 Moving a Transaction Log

You may want to move a transaction log if:

- You want to place the transaction log on a faster disk
- You want to redistribute the work load on your disks

How to Perform This Task

Caution

Follow all the steps carefully. Taking shortcuts can lead to data corruption.

1. Decide the location that you want to move the transaction log to, using the guidelines in Section 27.2.2. Remember that the disk must have enough contiguous space to hold the transaction log.
2. Log in to the node that the transaction log belongs to.
3. If you are in an OpenVMS Cluster, make sure that the disk you want to move the transaction log to is mounted clusterwide.
4. Decide which directory you want to move the transaction log to. You may want to create a new directory for the transaction log.
5. Find out which directory the transaction log is in, using LMCP's SHOW LOG command:

```
SHOW LOG SYSTEM$node.LM$JOURNAL
```

where *node* is the name of the node that the transaction log belongs to.

6. Rename the transaction log:

```
RENAME dirspecSYSTEM$node.LM$JOURNAL dirspecSYSTEM$node.LM$OLD
```

where:

dirspec is the full specification of the directory containing the transaction log.

node is the name of the node that the transaction log belongs to.

7. Can you stop all the software that uses DECdtm services without shutting down any nodes?

Managing DECdtm Services 27.8 Moving a Transaction Log

Yes Close the transaction log as follows:

Step	Action
a.	Stop all the software that uses DECdtm services.
b.	Close the transaction log using LMCP's CLOSE LOG command: <pre>\$ RUN SYS\$SYSTEM:LMCP LMCP> CLOSE LOG</pre> <p>The CLOSE LOG command closes the transaction log and stops the DECdtm TP_SERVER process. The command fails if any software is using DECdtm services.</p>
c.	Did the CLOSE LOG command succeed? <div style="margin-left: 20px;"> <p>Yes Restart the TP_SERVER process: <pre>\$ @SYS\$STARTUP:DECDTM\$STARTUP.COM</pre></p> <p>No Wait for 30 seconds, then repeat steps 7b and 7c.</p> </div>

No Close the transaction log by rebooting the node. Log in to the node when it has rebooted.

8. Make sure that SYS\$JOURNAL points to the directory that you want to move the log to. If SYS\$JOURNAL does not point to this directory, redefine SYS\$JOURNAL:

```
DEFINE/SYSTEM/EXECUTIVE_MODE SYS$JOURNAL dirspec[,...]
```

where *dirspec* is the full specification of a directory containing one or more transaction logs. List all the directories that will contain transaction logs after you have moved the transaction log. You can list the directories in any order.

In an OpenVMS Cluster, use SYSMAN to redefine SYS\$JOURNAL clusterwide.

9. If you redefined SYS\$JOURNAL in step 8, edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the definition of SYS\$JOURNAL.

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

10. Move the transaction log, using LMCP's CONVERT LOG command:

```
CONVERT LOG old-dirspecSYSTEM$node.LM$OLD new-dirspecSYSTEM$node.LM$JOURNAL
```

where:

old-dirspec is the full specification of the directory that currently contains the transaction log.

node is the name of the node that the transaction log belongs to.

new-dirspec is the full specification of the directory that you are moving the transaction log to.

11. If you stopped the software that uses DECdtm services in step 7, restart the software.

Managing DECdtm Services

27.8 Moving a Transaction Log

- Delete the old transaction log:

```
DELETE dirspecSYSTEM$node.LM$OLD;
```

where:

dirspec is the full specification of the directory containing the old transaction log.

node is the name of the node that the transaction log belongs to.

Example

This example shows how to move BLUE's transaction log. BLUE is in an OpenVMS Cluster. The cluster members and the locations of their transaction logs are as follows:

Node	Directory Containing Log
BLUE	DISK\$LOG1:[LOGFILES]
RED	DISK\$LOG2:[LOGFILES]

Neither node has a node-specific version of SYLOGICALS.COM.

Decide where you want to move BLUE's transaction log to. In this example, assume that you want to move it to DISK\$LOG3:[LOGFILES].

Log in to node BLUE. Then mount the disk clusterwide, and create a new directory for the transaction log:

```
$ MOUNT/CLUSTER/SYSTEM DUA3: LOG3
$ CREATE/DIRECTORY DISK$LOG3:[LOGFILES]
```

Find out which directory BLUE's transaction log is in, then rename the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> SHOW LOG SYSTEM$BLUE.LM$JOURNAL

Directory of DISK$LOG1:[LOGFILES]
SYSTEM$BLUE.LM$JOURNAL;1

Total of 1 file.
LMCP> EXIT
$ RENAME DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL -
_ $ DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$OLD
```

Stop all software that uses DECdtm services. Then close the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CLOSE LOG
Transaction log closed, TP_SERVER process stopped
LMCP> EXIT
```

Restart the TP_SERVER process:

```
$ @SYS$STARTUP:DECMTM$STARTUP.COM
```

Redefine SYS\$JOURNAL:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEFINE/SYSTEM/EXECUTIVE MODE SYS$JOURNAL -
SYSMAN> DISK$LOG2:[LOGFILES], DISK$LOG3:[LOGFILES]
SYSMAN> EXIT
```

Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the SYS\$JOURNAL definition. Then move the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CONVERT LOG DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$OLD -
  LMCP> DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL
Log file DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL;1 created.
Log file DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$OLD converted.
LMCP> EXIT
```

Restart the software that uses DECdtm services. Then delete the old transaction log:

```
$ DELETE DISK$LOG1:[LOGFILES]SYSTEM$BLUE.LM$OLD;
```

27.9 Dismounting a Disk

Before you can dismount a disk, you must close any transaction logs on the disk. This section describes how to dismount a disk that has transaction logs.

How to Perform This Task

1. Find out which transaction logs are on the disk you want to dismount, using LMCP's SHOW LOG command:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> SHOW LOG
```

2. Stop all the software that uses DECdtm services, if you can do so without shutting down any nodes.

If you cannot stop the software, reboot one or more nodes in step 3.

3. For each transaction log on the disk:

- a. Log in to the node that the transaction log belongs to.

- b. Rename the transaction log:

```
RENAME dirspecSYSTEM$node.LM$JOURNAL dirspecSYSTEM$node.LM$TEMP
```

where:

dirspec is the full specification of the directory containing the transaction log.

node is the name of the node that the transaction log belongs to.

- c. Did you stop all the software that uses DECdtm services in step 2?

Managing DECdtm Services

27.9 Dismounting a Disk

Yes Close the transaction log as follows:

Step	Action
1)	<p>Close the transaction log using LMCP's CLOSE LOG command:</p> <pre>\$ RUN SYS\$SYSTEM:LMCP LMCP> CLOSE LOG</pre> <p>The CLOSE LOG command closes the transaction log, and stops the DECdtm TP_SERVER process. The command fails if any software is using DECdtm services.</p>
2)	<p>Did the CLOSE LOG command succeed?</p> <p>Yes Restart the TP_SERVER process:</p> <pre>\$ @SYS\$STARTUP:DECDTM\$STARTUP.COM</pre> <p>No Wait for 30 seconds, then repeat step 3c.</p>

No Close the transaction log by rebooting the node. When the node has rebooted, log in.

4. Dismount the disk. For instructions on how to dismount a disk, see Section 9.9.
5. When you want to mount the disk again, follow these steps:
 - a. Mount the disk. For instructions on how to mount a disk, see Section 9.5. If you are in a cluster, mount the disk clusterwide.
 - b. Rename each transaction log on the disk:


```
RENAME dirspecSYSTEM$node.LM$TEMP dirspecSYSTEM$node.LM$JOURNAL
```

where:

dirspec is the full specification of the directory containing the transaction log.

node is the name of the node that the transaction log belongs to.
 - c. If you stopped the software that uses DECdtm services, restart the software.

Example

This example shows how to dismount the disk DISK\$LOG3.

Find out which transaction logs are on the disk:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> SHOW LOG
.
.
.
Directory of DISK$LOG3:[LOGFILES]
SYSTEM$BLUE.LM$JOURNAL;1)
```

The only transaction log on DISK\$LOG3 is node BLUE's transaction log.

Stop all the software that uses DECdtm services.

Log in to node BLUE. Then rename the transaction log:

```
$ RENAME DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL -
_ $ DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$TEMP
```

Close the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CLOSE LOG
Transaction log closed, TP_SERVER process stopped
LMCP> EXIT
```

Restart the TP_SERVER process:

```
$ @SYS$STARTUP:DECDTM$STARTUP.COM
```

Dismount the disk:

```
$ DISMOUNT/CLUSTER DISK$LOG3:
```

When you want to mount the disk again, mount it clusterwide:

```
$ MOUNT/CLUSTER/SYSTEM DUA3: LOG3
```

Rename BLUE's transaction log:

```
$ RENAME DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$TEMP -
_$ DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL
```

Restart the software that uses DECdtm services.

27.10 Adding a Node

For every node you add to an OpenVMS Cluster, you must create a new transaction log. This section describes how to create a transaction log for a new node.

How to Perform This Task

Before you perform this task, the new node must be configured into the cluster. For instructions on how to configure a node into a cluster, refer to *OpenVMS Cluster Systems*.

1. Decide the size and location of the new node's transaction log, using the guidelines in Section 27.2. Remember that the disk must have enough contiguous space to hold the log.
2. Make sure that the disk on which you want to create the transaction log is mounted clusterwide.
3. Decide which directory you want to create the new transaction log in. You may want to create a new directory for the transaction log.
4. Make sure that SYS\$JOURNAL points to the directory in which you want to create the new node's transaction log. If SYS\$JOURNAL does not point to this directory, use SYSMAN to redefine SYS\$JOURNAL clusterwide:

```
DO DEFINE/SYSTEM/EXECUTIVE_MODE SYS$JOURNAL dirspec[...]
```

where *dirspec* is the full specification of a directory containing one or more transaction logs. List all the directories that contain transaction logs, including the directory in which you want to create the new node's transaction log. You can list the directories in any order.

5. If you redefined SYS\$JOURNAL in step 4, edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the SYS\$JOURNAL definition.

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

Managing DECdtm Services

27.10 Adding a Node

6. Create the transaction log, using LMCP's CREATE LOG command:

```
CREATE LOG [/SIZE=size] dirspecSYSTEM$node.LM$JOURNAL
```

where:

size is the size of the transaction log in blocks. By default, the size of the transaction log is 4000 blocks.

dirspec is the full specification of the directory in which you want to create the transaction log.

node is the name of the new node.

Example

This example shows how to create a transaction log for a new node, whose SCSNODE name is WHITE.

In this example, the cluster members and the locations of their transaction logs are as follows:

Node	Directory Containing Log
BLUE	DISK\$LOG3:[LOGFILES]
RED	DISK\$LOG2:[LOGFILES]

Neither node has a node-specific version of SYLOGICALS.COM.

Decide the size and location of WHITE's transaction log:

Node	Size of Log (in Blocks)	Disk
WHITE	5000	DUA4

Mount the disk DUA4 clusterwide:

```
$ MOUNT/CLUSTER/SYSTEM DUA4: LOG4
```

Create a directory for the transaction log:

```
$ CREATE/DIRECTORY DISK$LOG4:[LOGFILES]
```

Redefine SYS\$JOURNAL:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEFINE/SYSTEM/EXECUTIVE MODE SYS$JOURNAL -
SYSMAN> DISK$LOG2:[LOGFILES], DISK$LOG3:[LOGFILES], DISK$LOG4:[LOGFILES]
SYSMAN> EXIT
```

Edit the SYS\$STARTUP:SYLOGICALS command procedure to update the SYS\$JOURNAL definition. Then create the transaction log:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CREATE LOG/SIZE=5000 DISK$LOG4:[LOGFILES]SYSTEM$WHITE.LM$JOURNAL
LMCP> EXIT
```

27.11 Removing a Node

This section describes how to remove a node if you are using DECdtm services.

How to Perform This Task

If you have a standalone machine, perform steps 1 to 8 only.

Caution

Follow all the steps carefully. Taking shortcuts can lead to data corruption.

1. Log in to the node that you want to remove.
2. Stop all the software that uses DECdtm services.
3. Find out whether the node's transaction log contains any active transactions, using LMCP's DUMP/ACTIVE command:

```
DUMP/ACTIVE SYSTEM$node.LM$JOURNAL
```

where *node* is the name of the node that you want to remove.
This command displays details of all the active transactions. The last line gives the total number of active transactions.
4. If the transaction log contains active transactions, follow these steps:
 - a. Run recovery procedures for all software that uses DECdtm services.
 - b. Find out if the node's transaction log still contains active transactions, using LMCP's DUMP/ACTIVE command.
 - c. If the transaction log still contains active transactions, contact your Compaq support representative.
5. Redefine SYS\$JOURNAL to exclude the directory that contains the transaction log of the node you want to remove, unless the directory contains other transaction logs.

```
DEFINE/SYSTEM/EXECUTIVE_MODE SYS$JOURNAL dirspec[...]
```

where *dirspec* is the full specification of a directory containing one or more transaction logs. List all the directories that contain any transaction logs other than the transaction log of the node you are removing. You can list the directories in any order.
In a cluster, use SYSMAN to redefine SYS\$JOURNAL clusterwide.
6. If you redefined SYS\$JOURNAL in step 5, edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the SYS\$JOURNAL definition.
If you created node-specific versions of SYLOGICALS.COM, edit all the versions.
7. Archive the transaction log.
8. Shut down the node.
9. Restart the software that uses DECdtm services.
10. Reconfigure the cluster to remove the node.

For information about how to reconfigure a cluster, refer to *OpenVMS Cluster Systems*.

Managing DECdtm Services

27.11 Removing a Node

Example

This example shows how to remove the node BLUE. In this example, the cluster members and the locations of their transaction logs are as follows:

Node	Directory Containing Log
BLUE	DISK\$LOG3:[LOGFILES]
RED	DISK\$LOG2:[LOGFILES]
WHITE	DISK\$LOG4:[LOGFILES]

None of the nodes has a node-specific version of the SYLOGICALS.COM command procedure.

Log in to node BLUE.

Stop all the software that uses DECdtm services. Then find out if BLUE's transaction log contains any active transactions:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> DUMP/ACTIVE SYSTEM$BLUE.LM$JOURNAL

Dump of log file DISK$LOG3:[LOGFILES]SYSTEM$BLUE.LM$JOURNAL
.
.
.
Total of 0 transactions active, 0 prepared and 0 committed.
LMCP> EXIT
```

Redefine SYS\$JOURNAL:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEFINE/SYSTEM/EXECUTIVE MODE SYS$JOURNAL -
SYSMAN> DISK$LOG2:[LOGFILES], DISK$LOG4:[LOGFILES]
SYSMAN> EXIT
```

Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to update the SYS\$JOURNAL definition.

Archive BLUE's transaction log. Then shut down node BLUE:

```
$ @SYS$SYSTEM:SHUTDOWN.COM
.
.
.
Should an automatic system reboot be performed [NO]? NO
```

Restart the software that uses DECdtm services. Then reconfigure the cluster:

```
$ @SYS$STARTUP:CLUSTER_CONFIG.COM

Cluster Configuration Procedure

1. ADD a node to a cluster.
2. REMOVE a node from the cluster.
3. CHANGE a cluster member's characteristics.
4. CREATE a duplicate system disk for BLUE.

Enter choice [1]: 2
.
.
.
Updating network database...
The configuration procedure has completed successfully.
```

27.12 Disabling DECdtm Services

By default, DECdtm services start automatically when you boot the computer. The DECdtm process, TP_SERVER, then checks for a transaction log, and continues checking until it finds one.

Disable DECdtm services if you do not use, and do not plan to use, any software that uses DECdtm services. This saves memory and CPU time.

In an OpenVMS Cluster, disable DECdtm services on all the nodes in the cluster.

How to Perform This Task

1. For each node:
 - a. Log in to the node.
 - b. Stop the TP_SERVER process using LMCP's CLOSE LOG command:

```
$ RUN SYS$SYSTEM:LMCP
LMCP> CLOSE LOG
```

The CLOSE LOG command stops the TP_SERVER process, providing no software is using DECdtm services.

If the CLOSE LOG command fails, do not continue this task.

If you have already stopped the TP_SERVER process on other nodes in a cluster system, restart the process using the SYS\$STARTUP:DECDTM\$STARTUP.COM command procedure.

2. Add the following line to the SYS\$MANAGER:SYLOGICALS.COM command procedure:

```
$ !
$ DEFINE/SYSTEM/EXECUTIVE_MODE SYS$DECDTM_INHIBIT yes
$ !
```

If you created node-specific versions of SYLOGICALS.COM, edit all the versions.

This stops the TP_SERVER process being created the next time you boot the system.

27.13 Enabling DECdtm Services

Enable DECdtm services only if you have previously disabled them and you now want to run software that uses DECdtm services.

How to Perform This Task

1. Deassign the logical name SYS\$DECDTM_INHIBIT:

```
$ DEASSIGN/SYSTEM/EXECUTIVE_MODE SYS$DECDTM_INHIBIT
```

In an OpenVMS Cluster, use SYSMAN to deassign SYS\$DECDTM_INHIBIT clusterwide.

2. Start up the DECdtm services process, TP_SERVER:

```
$ @SYS$STARTUP:DECDTM$STARTUP.COM
```

In an OpenVMS Cluster, use SYSMAN to start up the TP_SERVER process clusterwide.

Managing DECdtm Services

27.13 Enabling DECdtm Services

3. Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to delete the SYS\$DECDTM_INHIBIT definition. This ensures that DECdtm services start automatically when you boot the system.

Example

This example shows how to enable DECdtm services in a cluster environment.

Deassign SYS\$DECDTM_INHIBIT, then start up the TP_SERVER process.

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> SET ENVIRONMENT/CLUSTER
SYSMAN> DO DEASSIGN/SYSTEM/EXECUTIVE_MODE SYS$DECDTM_INHIBIT
SYSMAN> DO @SYS$STARTUP.DECDTM$STARTUP.COM
SYSMAN> EXIT
```

Edit the SYS\$MANAGER:SYLOGICALS.COM command procedure to delete the SYS\$DECDTM_INHIBIT definition.

27.14 Using the XA Gateway (Alpha Only)

DECdtm/XA provides support for coordinating and managing transactions that use XA. By using an XA Gateway, DECdtm/XA can join other Resource Managers (RM) in transactions that are managed by another Transaction Manager (TM). This section describes how to configure and use DECdtm XA Gateway support.

Note

In this chapter, the term **XA Specification** refers to *Distributed Transaction Processing: The XA Specification*.

To use DECdtm/XA and ensure proper startup and shutdown of DECdtm/XA services, the following command files must be invoked:

- SYS\$STARTUP:DDTM\$XA_STARTUP.COM
- SYS\$STARTUP:DDTM\$XA_SHUTDOWN.COM

Add the command @SYS\$STARTUP:DDTM\$XA_STARTUP.COM to the startup database or to the command file SYS\$MANAGER:SYSTARTUP_VMS.COM.

Add the command @SYS\$STARTUP:DDTM\$XA_SHUTDOWN.COM to the command file SYS\$MANAGER:SYSHUTDWN.COM.

Perform the following steps to verify that DECdtm XA services are operating properly:

1. Use the XGCP utility to create a gateway log file with the same name as the local OpenVMS node. See Section 27.14.1 and the *OpenVMS System Management Utilities Reference Manual* for details.
2. Run SYS\$TEST:DECDTM_XG_IVP.EXE.
3. Use the XGCP utility to stop and restart the gateway server. This step is essential if you choose to configure the gateway with a name different than that of the local OpenVMS node. For more information on the XGCP utility, see the *OpenVMS System Management Utilities Reference Manual*.

27.14.1 Gateway Configuration

The XA Gateway is configured into each transaction processing (TP) process as an XA-compliant resource manager. The XA Gateway handles XA calls from the XA transaction manager (TM) and maps them into calls to DECdtm system services. This allows DECdtm to send the appropriate events to any DECdtm compliant Resource Manager (RM) used in a TP process.

The operation of the XA Gateway is transparent to the RM; DECdtm RMs do not need any modification to be used with the XA Gateway.

The XA Gateway uses a log file to record the mapping between XA transactions and DECdtm transactions. The log file is managed by the gateway server process DDTM\$XG_SERVER.

Create the gateway log file with the XGCP utility (see the *OpenVMS System Management Utilities Reference Manual*). The size of the gateway log file depends on the number of concurrently active transactions. Up to 600 bytes are required for each active transaction, depending on the size of the transaction ID (TID) used by the XA TM. The gateway log file expands automatically when required.

The gateway log file resides in the directory specified by the logical name SYS\$JOURNAL and has a name of the form SYSTEM\$name.DDTM\$XG_JOURNAL. For optimum performance, move each gateway log file and each DECdtm log file to a separate physical device, and define SYS\$JOURNAL as a search list for the set of physical devices.

The XA Gateway requires an association on each OpenVMS Cluster node between an XA transaction manager and the XA Gateway log file. This association is managed by specifying a gateway name as follows:

- Use the XGCP utility to create a gateway log file with the gateway name (see the *OpenVMS System Management Utilities Reference Manual*).
- The gateway name is specified in the xa_open information string when the Gateway RM is configured in applications run under the control of an XA TM. (XA RM configuration is described in *OpenVMS Programming Concepts Manual, Volume II*.)
- The first XA application run by the XA TM binds the gateway name to the local node of the OpenVMS Cluster. The gateway name remains bound to that node until the gateway server is stopped.

All XA applications that run on the local node must be configured with the same gateway name. XA applications using the same name cannot run on other OpenVMS Cluster nodes. Therefore, you normally define one gateway name and create one gateway log file for each node of an OpenVMS Cluster.

You can change the association of a gateway name and bind the gateway name to a different OpenVMS Cluster node, provided that the node can access the gateway log file. To change the association of a gateway name, perform the following steps:

1. Stop all XA applications on the original node.
2. Use the XGCP utility to stop the gateway server on the original node.
3. Stop all XA applications on the new node.
4. Stop the gateway server on the new node and then restart the gateway server.

Managing DECdtm Services

27.14 Using the XA Gateway (Alpha Only)

5. Run the original XA applications on the new node.

Note

You must take care to protect against the loss of a gateway log file, perhaps by shadowing the device on which it resides. If you create a new log file, or if you use an out-of-date log file, transactions that were originally recorded as committed may be incorrectly rolled back. This can cause databases to become inconsistent with each other, or inconsistent with reports given to other systems or users.

Gateway log files are not large and it is better never to delete them. If you do delete an unwanted gateway log file, first use the DECdtm XGCP utility to verify that the gateway is not still a participant in any prepared transactions. The gateway participant name is DDTM\$XG/*name*.

The gateway server uses the following system logical names:

- **SYS\$DECDTM_XG_REQS**
Number of concurrent requests processed by the server, in the range 100 to 100,000. This determines the size of the global section that DDTM\$XG uses for communication with the server, and the quotas required by the server. The parameter is specified by defining the logical name SYS\$DECDTM_XG_REQS. Changes to the parameter do not take effect until after the server and all client processes have been stopped.
If the value of this parameter is exceeded during operation, client requests are blocked instead of being processed in parallel.
- **SYS\$DECDTM_XA_TRANS**
Estimated number of concurrent XA transactions, in the range 1000 to 1,000,000. This determines the size of indexing tables used internally in the server. The parameter is specified by defining the logical name SYS\$DECDTM_XA_TRANS. Changes to this parameter do not take effect until after the server has been stopped.
If the value of this parameter is exceeded during operation, server CPU use will increase. However, the effect is unlikely to be noticeable until the value of this parameter is exceeded by a factor of 10 or more.

Managing Special Processing Environments

The OpenVMS operating system supports the following special environments:

- Symmetric multiprocessing
- Vector processing (available only on certain CPU models)

This chapter describes how to set up and manage these special processing environments.

Information Provided in This Chapter

This chapter describes the following tasks:

Task	Section
Creating a multiprocessing environment	Section 28.2.1
Monitoring a multiprocessing environment	Section 28.2.2
†Loading the vector processing support code	Section 28.4.1
†Configuring a vector processing system	Section 28.4.2
†Managing vector processes	Section 28.4.3
†Restricting access to the vector processor with ACLs	Section 28.4.4
†Obtaining information about a vector processing system	Section 28.4.5
†Loading the VAX Vector Instruction Emulation facility (VVIEF)	Section 28.4.6
†VAX specific	

This chapter explains the following concepts:

Concept	Section
Symmetric multiprocessing	Section 28.1
Primary and secondary processors	Section 28.1.1
Available and active sets	Section 28.1.2
Vector processing	Section 28.3
†VAX support for vector processing	Section 28.3.1
†The VAX Vector Instruction Emulation facility (VVIEF)	Section 28.3.2
†VAX specific	

Managing Special Processing Environments

28.1 Understanding Multiprocessing

28.1 Understanding Multiprocessing

A multiprocessing system consists of two or more CPUs that address a common pool of memory and that are capable of executing instructions simultaneously.

The OpenVMS operating system supports a tightly coupled, symmetric multiprocessing (SMP) system. In a tightly coupled SMP system, all processors execute a single copy of the operating system and have equal access to all operating system code and system resources. OpenVMS SMP dynamically selects the CPU where a process will run based on process priority.

A multiprocessing system can function as an isolated entity, a node in a network, or a member of an OpenVMS Cluster environment. Multiprocessing and uniprocessing systems run the same operating system, although multiprocessing can be enabled only on selected VAX and Alpha processors. All processors in a multiprocessing environment must be at the same hardware and firmware level to guarantee that a given processor is capable of resuming the execution thread of a process that had been executing previously on another processor in the system.

28.1.1 Primary and Secondary Processors

In a multiprocessing system, one processor has the responsibility of starting other processors in the system. The **primary processor** is that processor in the system that is either logically or physically attached to the console device. As such, it is the processor that is the target of the console commands that boot the multiprocessing system. In this role, only the primary processor performs the initialization activities that define the operating system environment and prepare memory for the entire system. In addition, the primary processor serves as the system timekeeper, maintaining the system time and monitoring the timer queue for the expiration of its elements. In this sense, all processors in a multiprocessing system that do *not* have these responsibilities are known as **secondary processors**.

28.1.2 Available and Active Sets

An **available set** is made up of the processors that have passed the system's power-on hardware diagnostics and may or may not be actively involved in the system. Together, the primary and the secondary processors comprise the multiprocessing system's available set.

The **active set** is the subset of the VAX or Alpha system's processors that have passed power-on diagnostics and are actively participating in system operations. The operating system identifies each processor in these sets by its CPU ID, a value prevalent in the syntax and displays of certain DCL and utility commands.

28.1.3 Processor Capabilities

The processors in a multiprocessing system offer certain capabilities to the processes executing in the system. The following capabilities are supported:

- Primary
- Quorum
- Run
- Vector (VAX Only)

In addition, mechanisms exist to add and subtract other capabilities.

The Run capability affects CPU starting and stopping operations.

28.2 Managing SMP Environments

Managing symmetric multiprocessing systems (SMP) involves creating and monitoring a multiprocessing environment.

28.2.1 Creating a Multiprocessing Environment

You can control the membership and character of a multiprocessing system at boot time by setting system parameters designed for these purposes. Among the system parameters that manage a multiprocessing system are the following parameters:

Parameter	Function
MULTIPROCESSING	Determines which synchronization image is loaded into the operating system at boot time
SMP_CPUS	Determines which processors are brought into the multiprocessing environment from the available set at boot time

For more information about these and other system parameters, refer to the *OpenVMS System Management Utilities Reference Manual*.

You can add an available processor to the active set at boot time, or you can add it later using the DCL command `START/CPU`. The DCL command `STOP/CPU` removes a processor from the active set.

SMP Extension License

On Alpha systems, you must register the SMP Extension License if you have an SMP system. This license upgrades the Operating System Base License and all Interactive User licenses to the matching multiprocessing level of your system.

Because the SMP Extension License grants all the rights the existing Base and User licenses provide at the uniprocessing level, reinstalling those licenses when you upgrade to a multiprocessing system is unnecessary. When your system is upgraded to a new multiprocessing level, add an SMP Extension License to your existing license.

28.2.2 Monitoring a Multiprocessing Environment

Several operating system features provide special information about the character, capabilities, and status of a multiprocessor system. They include the DCL command `SHOW CPU` and the Monitor utility.

Obtaining Information About a Multiprocessor Configuration

The `SHOW CPU` command displays three levels of information describing the configuration and status of a multiprocessing system:

Level	Command Example	Display Contents
Summary	<code>SHOW CPU</code>	Indicates which processor is primary, which processors are configured, and which processors are active; displays the minimum revision levels for processors in the system and the setting of the <code>MULTIPROCESSING</code> system parameter; and indicates whether multiprocessing is enabled.

Managing Special Processing Environments

28.2 Managing SMP Environments

Level	Command Example	Display Contents
Brief	SHOW CPU/BRIEF	Produces information from the summary display; lists the current CPU state and the current process (if any) for each configured processor.
Full	SHOW CPU/FULL	Produces information from the summary display and displays additional information; lists the current CPU state, current process (if any), revision levels, and capabilities for each configured processor; indicates which processes can be executed only on certain processors.

For more information about the DCL commands relating to SMP, refer to the *OpenVMS DCL Dictionary*; for information about the Monitor utility, refer to the MONITOR section in the *OpenVMS System Management Utilities Reference Manual*.

28.3 Understanding Vector Processing

A single data item, having one value, is known as a **scalar**. A group of related scalar values, or elements, all of the same data type, is known as a **vector**.

Traditional (scalar) computers operate only on scalar values, and must process vector elements sequentially. Vector computers, on the other hand, recognize vectors as native data structures and can operate on an entire vector with a single vector instruction. Because this type of processing involves the concurrent execution of multiple arithmetic or logical operations, a vector computer can routinely process a vector four to five times faster than a traditional computer can using only scalar instructions.

Vector processors gain a further speed advantage over scalar processors by their use of special hardware techniques designed for the fast processing of streams of data. These techniques include data pipelining, chaining, and other forms of hardware parallelism in memory and in arithmetic and logical functional units. Pipelined functional units allow the vector processor to overlap the execution of successive computations with previous computations.

28.3.1 VAX Support for Vector Processing (VAX Only)

The VAX vector architecture includes sixteen 64-bit vector registers (V0 through V15), each containing 64 elements; vector control registers, including the vector count register (VCR), vector length register (VLR), and vector mask register (VMR); vector functional units; and a set of vector instructions. VAX vector instructions transfer data between the vector registers and memory, perform integer and floating-point arithmetic, and execute processor control functions. A more detailed description of the VAX vector architecture, vector registers, and vector instructions appears in the *VAX MACRO and Instruction Set Reference Manual*.

Those VAX systems that comply with the VAX vector architecture are known as **vector-capable systems**.

A VAX vector processing system configuration includes one or more integrated scalar-vector processor pairs, or **vector-present processors**. Such a configuration can be symmetric, including a vector coprocessor for each scalar, or asymmetric, incorporating additional scalar-only processors. Depending upon the model of the VAX vector processing system, the scalar and vector CPUs of vector-present processors can be either a single, integral physical module or separate, physically independent modules. In either case the scalar and vector

Managing Special Processing Environments

28.3 Understanding Vector Processing

CPUs are logically integrated, sharing the same memory and transferring data over a dedicated, high-speed internal path.

Like VAX scalar processing systems, a VAX vector processing system can participate as a member of a VAXcluster or a node in a network, or be run as a standalone system.

28.3.2 VAX Vector Instruction Emulation Facility (VAX Only)

The VAX Vector Instruction Emulation Facility (VVIEF) is a standard feature of the OpenVMS operating system that allows vectorized applications to be written and debugged in a VAX system in which vector processors are not available. VVIEF emulates the VAX vector processing environment, including the nonprivileged VAX vector instructions and the vector system services. Use of VVIEF is restricted to user mode code.

VVIEF is strictly a program development tool, and *not* a run-time replacement for vector hardware. Vectorizing applications to run under VVIEF offers no performance benefit; vectorized applications running under VVIEF execute more slowly than their scalar counterparts.

The operating system supplies the VVIEF bootstrap code as an executive loadable image. Note that, in the presence of OpenVMS vector support code, VVIEF remains inactive. Although it is possible to prevent the loading of vector support code in a vector-present system (see Section 28.4.1) and activate VVIEF, there are few benefits.

See Section 28.4.6 for additional information about loading and unloading VVIEF.

28.4 Managing the Vector Processing Environment (VAX Only)

The following sections describe tasks for managing a vector processing system.

28.4.1 Loading the Vector Processing Support Code (VAX Only)

By default, in a VAX vector processing system, the system automatically loads the vector processing support code at boot time. You can override the default behavior by setting the static system parameter VECTOR_PROC as described in Table 28–1.

Table 28–1 Settings of VECTOR_PROC System Parameter (VAX Only)

Value	Result
0	Do not load the vector processing support code, regardless of the system configuration.
1	Load the vector processing support code if at least one vector-present processor exists. This is the default value.
2	Load the vector processing support code if the system is vector-capable. This setting is most useful for a system in which processors have separate power supplies. With this setting, you can reconfigure a vector processor into the system without rebooting the operating system.

Managing Special Processing Environments

28.4 Managing the Vector Processing Environment (VAX Only)

28.4.2 Configuring a Vector Processing System (VAX Only)

You can add a vector-present processor to or remove it from a multiprocessing configuration at boot time by using the system parameter `SMP_CPUS`, or at run time by using the DCL commands `START/CPU` and `STOP/CPU`. Note that the operating system treats the scalar and vector CPU components of a vector-present processor as a single processor, starting them and stopping them together.

At boot time, the setting of the system parameter `SMP_CPUS` identifies which secondary processors in a multiprocessing system are to be configured, including those processors that are vector present. (The operating system always configures the primary processor.) The default value of `-1` boots all available processors, scalar and vector-present alike, into the configuration. (Refer to the *OpenVMS System Management Utilities Reference Manual* for additional information about this parameter.) Note that, prior to starting a vector-present processor, you should ensure that the vector processing support code (see Section 28.4.1) is loaded at boot time. Otherwise, processes will be able to use only the scalar CPU component of the vector-present processor.

To bring secondary processors into a running multiprocessing system, use the DCL command `START/CPU`. To remove secondary processors from the system, use the `STOP/CPU` commands. Again, you must ensure that the vector processing support code has been loaded at boot time for the vector CPU component of vector-present processors started in this way to be used.

Note, however, that a `STOP/CPU` command fails and generates a message if it would result in the removal of a vector-present processor that is the sole provider of the vector capability for currently active vector consumers. In extreme cases, such as the removal of a processor for repair, you can override this behavior by issuing the command `STOP/CPU/OVERRIDE`. This command stops the processor, despite stranding processes.

When a `STOP/CPU/OVERRIDE` command is issued for a vector-present processor, or when a vector-present processor fails, the operating system puts all stranded vector consumers into a CPU-capability-wait (`RSN$_CPUCAP`) state until a vector-present processor is returned to the configuration. To any other process that subsequently issue a vector instruction (including a marginal vector consumer), the system returns a “requested CPU not active” message (`CPUNOTACT`).

Refer to the *OpenVMS DCL Dictionary* for additional information about the `START/CPU` and `STOP/CPU` commands.

28.4.3 Managing Vector Processes (VAX Only)

The operating system scheduling algorithms automatically distribute vector and scalar processing resources among vector consumers, marginal vector consumers, and scalar consumers. However, VAX vector processing configurations vary in two important ways:

- The amount of vector processing activity the configuration must accommodate
- The number of vector-present processors that are available in the configuration to service vector processing needs

In a configuration that has more vector consumers in a system than scalar-vector processor pairs to service them, vector consumers share vector-present processors according to process priority. At a given priority, the system schedules vector consumers on a vector-present processor in a round-robin fashion. Each time the system must schedule a new vector consumer on a vector-present processor,

Managing Special Processing Environments

28.4 Managing the Vector Processing Environment (VAX Only)

it must save the vector context of the current vector consumer in memory and restore the vector context of the new vector consumer from memory. When such “slow” vector context switches occur too frequently, a significant portion of the processing time is spent on vector context switches relative to actual computation.

Systems that have heavy vector processing needs should be adequately configured to accommodate those needs. However, some mechanisms are available for tuning the performance of an existing configuration.

28.4.3.1 Adjusting System Resources and Process Quotas (VAX Only)

Systems in which several vector consumers are active simultaneously may experience increased paging activity as processes share the available memory. To reduce process paging, you may need to use the Authorize utility (AUTHORIZE) to adjust the working set limits and quotas of the processes running vectorized applications. (Refer to the AUTHORIZE section of the *OpenVMS System Management Utilities Reference Manual* for additional information.) An increase of the process maximum working set size (system parameter WSMAX) may also be necessary. Additionally, a vectorized application may use the Lock Pages in Working Set system service (\$LKWSET) to enhance its own performance.

The system allots to each vector consumer 8KB of system nonpaged dynamic memory in which the operating system stores vector context information. Depending upon how many vector consumers may be active in the system simultaneously, you may need to adjust the system parameter NPAGEDYN. The DCL command SHOW MEMORY/POOL/FULL displays the current size of nonpaged pool in bytes.

To obtain optimal performance of a VAX vector processing system, you should take some care in setting up generic batch queues that avoid saturating the system’s vector resources. If a queue contains more active vectorized batch jobs than vector-present processors in the system, a significant portion of the processing time will be spent on vector context switches.

The recommended means for dispatching vectorized batch jobs to a VAX vector processing system is to set up a separate queue (for instance, VECTOR_BATCH) with a job limit equal to the number of vector-present processors in the system. When submitting vectorized batch jobs, users should be encouraged to submit them to this generic vector-processing batch queue.

28.4.3.2 Distributing Scalar and Vector Resources Among Processes (VAX Only)

As a vector consumer, a process must be scheduled only on a vector-present processor. If the image the process is executing issues only scalar instructions for a period of time, and it must share the scalar-vector processor pair with other vector consumers, its inability to run on an available scalar processor could hamper its performance and the overall performance of the system.

By default, the operating system assumes that if a vector consumer has not issued a vector instruction for a certain period of time, it is unlikely that it will issue a vector instruction in the near future. The system relinquishes this process’s need for the vector capability, classifying it as a marginal vector consumer.

In an asymmetric vector-processing configuration, detection of marginal vector consumers achieves the following desirable effects:

- Because a marginal vector consumer is eligible to run on a larger set of processors, its response time will improve.
- The scheduling of marginal vector consumers on scalar processors reduces the contention for vector-present processors.

Managing Special Processing Environments

28.4 Managing the Vector Processing Environment (VAX Only)

- Because vector consumers issuing vector instructions are more likely to be scheduled on vector-present processors, the vector CPU is more efficiently used.

Use the VECTOR_MARGIN system parameter to establish the interval of time at which the system checks the status of all vector consumers. The VECTOR_MARGIN parameter accepts an integer value between 1 and FFFFFFFF₁₆. This value represents a number of consecutive process quanta (as determined by the system parameter QUANTUM). If the process has not issued any vector instructions in the specified number of quanta, the system declares it a marginal vector consumer.

The default value of the VECTOR_MARGIN parameter is 200₁₀.

28.4.4 Restricting Access to the Vector Processor by Using ACLs (VAX Only)

A vector **capability** is a software abstract by which the operating system makes the services of the vector processor available to users. You can restrict the use of the vector processor to users holding a particular identifier by associating an access control list (ACL) with the vector capability object.

For example, a university might limit use of the vector processor to faculty and students in an image processing course, or a service bureau might charge users for access to the vector capability, time spent on the vector processor, or both.

Use the DCL command SET SECURITY/ACL in the following format to establish access control entries (ACEs) on a vector capability:

```
SET SECURITY /CLASS=CAPABILITY /ACL=(ace[,...]) VECTOR
```

The following DCL command displays the ACL on the vector capability:

```
$ SHOW SECURITY /CLASS=CAPABILITY VECTOR
```

Note that the ACL is on the vector capability, not on the use of any or all vector-present processors in the system. The operating system will still schedule processes without permission to use the vector capability on a vector-present processor. However, these processors will be able to use only the scalar CPU component of the processor, and cannot execute vector instructions. Likewise, because the ACL is on the vector capability and not on a vector-present processor, you cannot establish an ACL to force long-running jobs to a specific processor.

For additional information about the SET SECURITY and SHOW SECURITY commands, refer to the *OpenVMS DCL Dictionary*.

28.4.5 Obtaining Information About a Vector Processing System (VAX Only)

You can obtain information about the status of the vector processing system and the use of the system by individual processes through various means, including:

- The DCL lexical functions F\$GETJPI and F\$GETSYI
- The DCL command SHOW CPU
- The DCL commands SHOW PROCESS and LOGOUT/FULL
- The Accounting utility
- The Monitor utility

Managing Special Processing Environments

28.4 Managing the Vector Processing Environment (VAX Only)

28.4.5.1 DCL Lexical Functions F\$GETJPI and F\$GETSYI (VAX Only)

The DCL lexical function F\$GETJPI accepts the following items and returns the corresponding information regarding the vector status of a specified process:

Item	Return Type	Information Returned
FAST_VP_SWITCH	Integer	Number of times this process has issued a vector instruction that resulted in an inactive vector processor being enabled without the expense of a vector context switch
SLOW_VP_SWITCH	Integer	Number of times this process has issued a vector instruction that resulted in an inactive vector processor being enabled with a full vector context switch
VP_CONSUMER	Boolean	Flag indicating whether the process is a vector consumer
VP_CPUTIM	Integer	Total amount of time the process has accumulated as a vector consumer

The DCL lexical function F\$GETSYI accepts the following items and returns the corresponding information regarding the status of the vector processing system:

Item	Return Type	Information Returned
VECTOR_EMULATOR	Integer	Flag indicating the presence of the VAX Vector Instruction Emulation facility (VVIEF) in the system
VP_MASK	Integer	Mask indicating which processors in the system have vector coprocessor
VP_NUMBER	Integer	Number of vector processors in the system

Refer to the *OpenVMS DCL Dictionary* for additional information about the DCL lexicals F\$GETJPI and F\$GETSYI.

28.4.5.2 SHOW CPU/FULL Command (VAX Only)

The SHOW CPU/FULL command lists the capabilities of the specified CPU. Issue this command to determine the presence of the vector capability in the system prior to executing a STOP/CPU command.

Refer to the *OpenVMS DCL Dictionary* for additional information about the SHOW CPU command.

28.4.5.3 SHOW PROCESS and LOGOUT/FULL Commands (VAX Only)

If the target process has accrued any time as a vector consumer scheduled on a vector-present processor, the DCL commands SHOW PROCESS and LOGOUT/FULL display the elapsed vector CPU time and the charged vector CPU time, respectively.

To accumulate vector CPU time, a process must be a vector consumer (that is, require the system vector capability) and be scheduled on a vector-present processor. The operating system still charges the vector consumer vector CPU time, even if, when scheduled on the vector-present processor, it does not actually use the vector CPU. Note that, because scalar consumers and marginal vector consumers do not use the vector CPU, they do not accrue vector CPU time, even when scheduled on a vector-present processor.

Refer to the *OpenVMS DCL Dictionary* for additional information about the SHOW PROCESS and LOGOUT commands.

Managing Special Processing Environments

28.4 Managing the Vector Processing Environment (VAX Only)

28.4.6 Loading the VAX Vector Instruction Emulation Facility (VVIEF) (VAX Only)

The VAX Vector Instruction Emulation Facility (VVIEF) is a standard operating system feature that allows vectorized applications to be written and debugged in a VAX system in which vector processors are not available. VVIEF is intended strictly as a program development tool, and *not* as a run-time replacement for vector hardware. Vectorizing applications to run under VVIEF offers no performance benefit; vectorized applications running under VVIEF will execute more slowly than their scalar counterparts.

To cause the system to load VVIEF at the *next* system boot and at each subsequent system boot, invoke the command procedure SYS\$UPDATE:VVIEF\$INSTAL.COM. To unload VVIEF, invoke the command procedure SYS\$UPDATE:VVIEF\$DEINSTAL.COM and reboot the system.

You can determine the presence or absence of VVIEF in a system by issuing the following DCL commands:

```
$ X = F$GETSYI("VECTOR_EMULATOR")
$ SHOW SYMBOL X
X = 1   Hex = 00000001   Octal = 000000001
```

A return value of 1 indicates the presence of VVIEF; a value of 0 indicates its absence.

Note that, although VVIEF may be loaded into the system, in the presence of vector support code, it remains inactive. Although it is possible to prevent the loading of vector processing support code in a vector-present system (see Section 28.4.1) and activate VVIEF, there are few benefits. Should the only vector-present processor in the system fail, the execution of preempted vectorized applications will not resume under VVIEF.

Files-11 Disk Structure

This appendix explains disk terminology and disk concepts. It also describes reserved files, points out those files used by the Analyze/Disk_Structure utility (ANALYZE/DISK_STRUCTURE), and compares Files-11 On-Disk Structure (ODS) Level 1 with Files-11 ODS Levels 2 and 5.

A.1 Disk Concepts

This section defines terms related to both the physical and the logical organization of disks.

A.1.1 Logical Organization of a Disk

The smallest addressable unit of information on a disk is a **block**. Files-11 On-Disk Structures define a block to consist of 512 8-bit bytes. Blocks can be treated as units for transfer between a Files-11 disk volume and memory. Files-11 ODS, however, views a disk as an array of blocks, and is generally not concerned with individual blocks.

Blocks are logically grouped into **clusters**, which are the basic units by which disk space is allocated. You determine the number of blocks in a cluster when a given disk, known as a **volume**, is first prepared for use (initialized). Cluster sizes vary for different media types. The smaller cluster sizes in the range are usually more practical. In general, a disk with a relatively small number of blocks is given a smaller cluster size, while larger disks are given larger cluster sizes to minimize the overhead for disk space allocation.

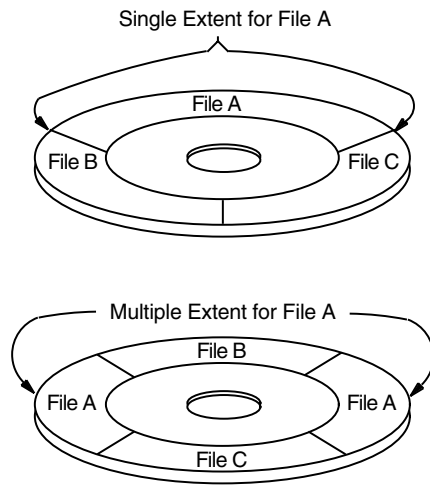
Contiguous clusters allocated to a particular file are called **extents**. An extent can contain all or part of a file. If enough contiguous area is available on the disk, the entire file is allocated as a single extent. Sometimes, however, not enough contiguous area is available to hold the entire file, or, when you create a file initially, you might not want to reserve the entire required amount of space. When the file is eventually extended, it is unlikely that the adjacent clusters will still be unallocated. If the adjacent clusters are already allocated to another file, the extension does not occur contiguously.

If a file is divided into two or more parts, each part is an extent. Thus, a file can consist of multiple extents located in separate areas on the disk, as shown in Figure A-1. Note that the file extensions are done automatically.

Files-11 Disk Structure

A.1 Disk Concepts

Figure A-1 File Extents



ZK-0738-GE

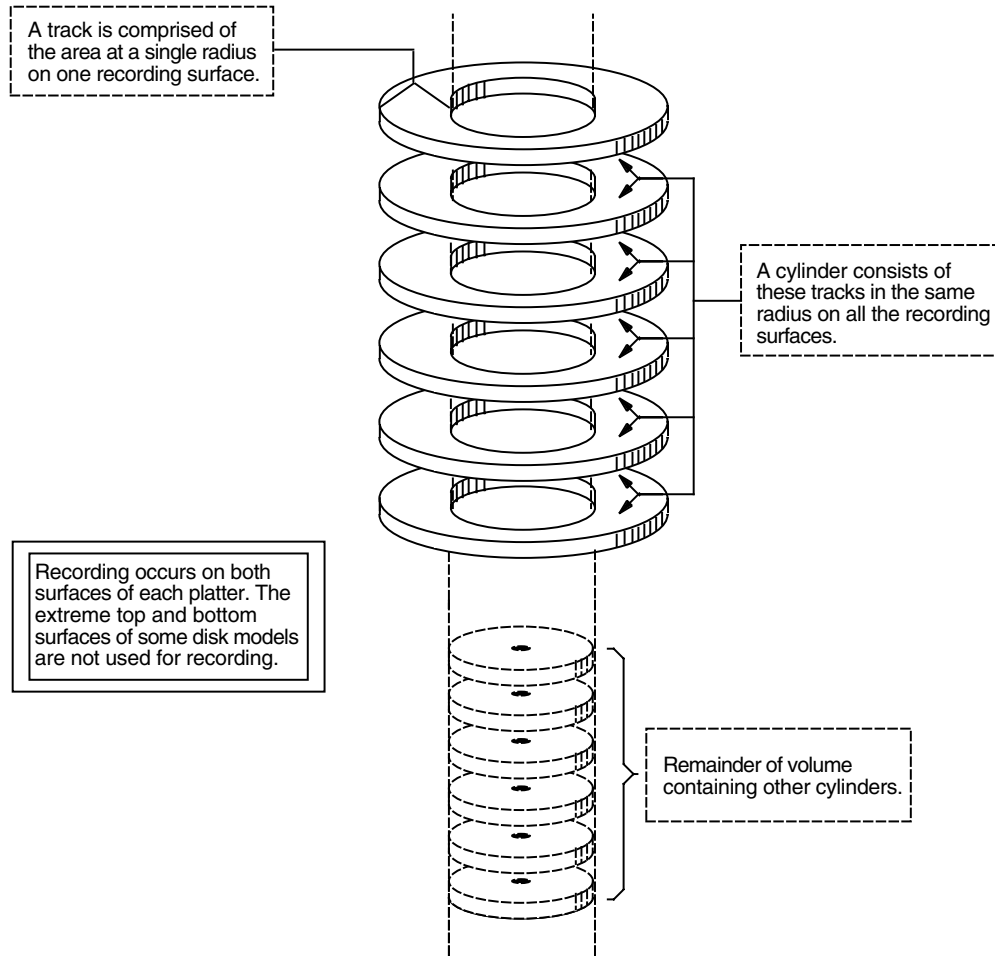
A.1.2 Physical Organization of a Disk

The smallest unit discernible to the Files-11 structure is the **sector**; for most Files-11 disks, a sector is equivalent to a block, which is 512 bytes. Other basic terms related to disks are **track** and **cylinder**. A track is the collection of sectors (or blocks, on Files-11 structures) at a single radius on one recording surface of a disk. It is accessible to a given read/write head position on the disk device. A cylinder consists of all tracks at the same radius on all recording surfaces of a disk.

Because access to any of the blocks in a given cylinder does not require any movement of the disk's read/write heads, it is generally advantageous to keep related data blocks in the same cylinder. For this reason, when choosing a cluster size for a large-capacity disk, you should usually select a cluster size that divides evenly into the cylinder size.

Figure A-2 is a graphic representation of disk tracks and cylinders.

Figure A-2 Tracks and Cylinders



ZK-0740-GE

A.2 Files-11 Structure

The Files-11 structure creates a set of nondeletable reserved files when a volume or volume set is initialized. These files control the organization of a Files-11 disk. A Files-11 structure resides on a volume, which is a physical medium such as a disk pack. A Files-11 volume is an ordered set of 512-byte blocks. The blocks are numbered consecutively from 0 to $n-1$; the value of $n-1$ is the size of the disk in blocks.

Files-11 Disk Structure

A.2 Files-11 Structure

A.2.1 File Identification (FID)

Each file on a Files-11 disk is identified by a unique, system-assigned file identification (FID) and can have a user-assigned alphanumeric name. The primary function of a Files-11 directory is to associate the user-assigned alphanumeric name of each file with the unique FID of the file. This association ensures that files present on a volume are retrievable by name.

The FID of a file consists of a set of three numbers. The first is the *file number* (NUM). The file system uses this number as an offset into the index file (reserved file INDEXF.SYS), which stores information for all files on a volume.

The second part of the FID is the *file sequence number* (SEQ), which represents the number of times a particular file number has been used. File numbers are allocated and deallocated as files are created and deleted. As a result, the file number alone cannot uniquely identify the file. By incrementing the sequence number each time a file number is used, the file system ensures that each file has a unique identification in INDEXF.SYS.

The third number in the FID is the *relative volume number* (RVN). This number indicates the volume (of a volume set) on which the file resides (ODS-2 only). If the volume set consists of a single volume, the RVN of all files on that volume is 1.

A.2.2 ODS Directory Hierarchies

The Files-11 ODS-2 structure is a multilevel directory hierarchy. The top level of the directory structure is the master file directory (MFD). The MFD of a volume is always named [000000]. The MFD contains all top-level directories, including itself, and reserved files.

A directory is a file that contains other files. A file contained in a directory can also be a directory and contain other files. By nesting directories, users can construct directory hierarchies up to nine levels deep (including the master file directory).

In a volume set, the MFD for all of the user directories on the volume set is located on relative volume 1. The entries of this MFD point to directories located on any volume in the set. These directories in turn point to files and subdirectories on any volume in the set. The MFD of any remaining volume in the set includes only the names of the reserved files for that volume.

On VAX systems, the Files-11 ODS-1 structure supports a two-level directory hierarchy. Each user identification code (UIC) is associated with a user file directory (UFD). Each UFD is included in the MFD of the volume.

A.3 Reserved Files

This section describes the reserved files that Files-11 uses. Note that all reserved files have constant FIDs.

This section also points out the files ANALYZE/DISK_STRUCTURE uses. ANALYZE/DISK_STRUCTURE makes an in-memory copy of what these files should look like and compares it with the current version. The utility reports and repairs (if you specify the /REPAIR qualifier) any discrepancies found during these comparisons.

Table A-1 shows the reserved files used by Files-11 Levels 1, 2, and 5, and files used by ANALYZE/DISK_STRUCTURE.

Table A-1 Reserved Files

Reserved File	File Name	†Structure Level 1	Structure Levels 2 and 5	ANALYZE/ DISK_ STRUCTURE
Index file	INDEXF.SYS;1	X	X	X
Storage bitmap file	BITMAP.SYS;1	X	X	X
Bad block file	BADBLK.SYS;1	X	X	
Master file directory	000000.DIR;1	X	X	X
Core image file	CORIMG.SYS;1	X	X	
Volume set list file	VOLSET.SYS;1		X	X
Continuation file	CONTIN.SYS;1		X	
Backup log file	BACKUP.SYS;1		X	
Pending bad block	BADLOG.SYS;1		X	
Quota file	QUOTA.SYS			X
Volume security profile	SECURITY.SYS		X	

†VAX specific

A.3.1 Index File, INDEXF.SYS

Every Files-11 volume has an index file, which is created when the volume is initialized. (You cannot use a disk as a Files-11 disk until it has been initialized with the INITIALIZE command.)

INDEXF.SYS is a large, extendable file made up of several sections. These sections provide the operating system with the information necessary to identify a Files-11 volume, initially access that volume, and locate all the files on that volume (including INDEXF.SYS itself).

Table A-2 shows the information that is in INDEXF.SYS. After the table are additional explanations of boot block, home block, and file headers.

Table A-2 Contents of Files-11 Index File

Term	Definition
Boot block	Virtual block 1 of the index file. The boot (or bootstrap) block is almost always mapped to the logical block 0 of the volume. If the volume is a system volume, the boot block contains a boot program that loads the operating system into memory. If the volume is not a system volume, the boot block contains a program that displays the message that the volume is not the system device but a device that contains users' files only.
Home block	Establishes the specific identity of the volume, providing such information as the volume name and protection, the maximum number of files allowed on the volume, and the volume ownership information. The home block is virtual block number 2 of the index file.

(continued on next page)

Files-11 Disk Structure

A.3 Reserved Files

Table A-2 (Cont.) Contents of Files-11 Index File

Term	Definition
Backup home block	A copy of the home block; permits the volume to be used even if the primary home block is destroyed.
Backup index file header	Permits data on the volume to be recovered if the index file header is corrupted; occupies virtual blocks $v * 3 + 1$ through $v * 4$, where v is the volume cluster factor.
Index file bitmap	Controls the allocation of file headers and thus the number of files on the volume; contains a bit for each file header allowed on the volume. If the value of a bit for a given file header is 0, a file can be created with this file header. If the value is 1, the file header is already in use.
File headers	Make up the bulk of the index file; contain all the information needed for gaining access to the file. Each file header describes one file on the volume. A file header contains information such as the owner UIC, protection code, creation date and time, and access control lists (ACLs); it also contains a list of extents that make up the file, describing where the file is logically located on the volume. Note that a file header can also be an extension header.
Alternate index file header	Permits recovery of data on the volume if the primary index file header becomes damaged.

A.3.1.1 Boot Block

Block 0 on a system disk is the **boot block**. It contains the location and size of the **primary bootstrap image**, which is used to boot the system. Certain processors, in order to boot, must read this boot block to obtain the location of the bootstrap image. For more details, see Section 4.7.

A.3.1.2 Home Block

The **home block** is normally the next block after the boot block; it identifies the disk as a Files-11 volume. If for some reason the home block cannot be read (physically unusable), an alternative block will be selected for use as the home block. This block provides specific information about the volume and default values for files on the volume. Items in the home block include the following ones:

- The volume name
- Information to locate the remainder of the index file
- The maximum number of files that can be present on the volume at any one time
- The user identification code (UIC) of the owner of the volume
- Volume protection information (specifies which users can read or write the entire volume)

Files-11 volumes contain several copies of the home block to ensure against accidental destruction of this information and the consequent loss of access to files on the volume.

A.3.1.3 File Headers

Most of the index file consists of **file headers**; each file header describes a portion of a file on the volume. File headers contain information such as the owner UIC, protection code, creation date and time, and access control lists (ACLs). Most importantly, the file header contains a list of extents that make up the file, describing where the file is logically located on the volume. If a file has a large number of extents, multiple file headers may be used to describe them. A file identifier number is associated with each file header.

When you create a file, you normally specify a file name to OpenVMS RMS, which assigns this name to the file on a Files–11 volume. OpenVMS RMS places the file name and file identifier associated with the newly created file into a directory, which contains an entry defining the location for each file. When you access the file, you supply the file name, which supplies a path to the file identifier through the directory entry. The file identifier, in turn, points to the location of the file header, which contains a listing of the extent or extents that locate the actual data.

Because they represent the current state of file storage on a volume, file headers are of particular interest to ANALYZE/DISK_STRUCTURE. Each file on a Files–11 disk (INDEXF.SYS included) is identified and located by a primary header (and extension headers, if required) in INDEXF.SYS.

Each fixed-length header contains both constant and variable-length data. This data is stored in one of the six areas shown in Table A–3.

Table A–3 Areas of Data in File Headers

Area of Data	Description
Header	This area contains the header identification, the file number and its sequence number, the protection code for the file, and offsets to the other file header areas.
Ident	This area contains the identification and accounting data for the file (for example, the name of the file, its creation date and time, and backup date and time).
Map	This area contains a list of retrieval pointers that map the virtual blocks of the file to the logical blocks of the volume. Each pointer describes one group of consecutively numbered logical blocks that is allocated to the file. Retrieval pointers are arranged in the order of the virtual blocks they represent.
Access control list	An optional area that contains ACL-related information.
Reserved	This area is reserved for use by special applications.
End checksum	The last two bytes of the file header contain a 16-bit additive checksum of the preceding 255 words of the file header. The checksum helps verify that the block is a valid file header.

A set of contiguous clusters is known as an **extent**. The size of an extent varies according to the number of contiguous clusters. For example, assume a file requires 1000 blocks of storage, and the file system finds a set of 800 contiguous blocks and a set of 200 contiguous blocks. The file would then be stored in two extents: one consisting of 800 blocks, the other of 200.

Files–11 Disk Structure

A.3 Reserved Files

The *primary header* of a file points to the first extent of that file and to as many extents as can be stored in the map area of the primary header. When the number of extents required to contain a file exceeds the map area available in the primary header, or the ACL is too large to fit in the primary header, the file is allocated an *extension header*. Extension headers contain all the constant data of the primary header, as well as the variable data (in the header map area and access control list) that specifies the locations of the extents to which the extension header points.

ANALYZE/DISK_STRUCTURE confirms the validity of a file by working its way down the list of primary and extension headers of the file. During this process, ANALYZE/DISK_STRUCTURE checks the validity of the file header, the chain of pointers to all extension headers, the retrieval pointers in all headers, and the attributes of the file.

A.3.2 Storage Bitmap File, BITMAP.SYS

The storage bitmap file is a contiguous file that the file system uses to keep track of the available space on a volume. This file contains a storage control block (SCB), which consists of summary information intended to optimize the Files–11 space allocation, and the bitmap itself, which lists the availability of individual blocks.

The SCB contains summary information about the volume (cluster factor, volume size, blocking factor, and so forth). Each bit in the bitmap represents an allocatable cluster on the volume. If a bit is set, the corresponding cluster is available for use. If a bit is clear, the cluster is not available.

During normal operation, the operating system moves portions of the bitmap in and out of cache memory. The state of each bit in memory is altered as clusters are allocated and deallocated. BITMAP.SYS is updated when the portion of the bitmap in cache is swapped back to disk. Since a portion of the bitmap is always in cache, BITMAP.SYS never reflects the current state of allocated clusters on a disk (unless the disk is dismounted or write-locked).

One of the functions of ANALYZE/DISK_STRUCTURE is to build a current version of BITMAP.SYS from data extracted from INDEXF.SYS, so that BITMAP.SYS accurately reflects the status of free clusters on the disk.

A.3.3 Bad Block File, BADBLK.SYS

The bad block file contains all the bad blocks on the volume. The system detects bad disk blocks dynamically and prevents their reuse once the files to which they are allocated have been deleted.

A.3.4 Master File Directory

The MFD is a file that contains reserved files that control the Files-11 volume directory structure. The MFD lists the known files, in addition to any files or directories that the user enters. The master file directory is itself one of the files (000000.DIR;1) listed in the MFD.

Usually, however, the MFD is used to list the reserved files and users' file directories; users seldom enter files into the MFD, even on private volumes. In fact, on a private volume, it is most convenient for users to create a directory that has the same name as their default directory on a system disk. For an explanation of users' file directories and file specifications, refer to the *OpenVMS User's Manual*.

When the Backup utility (BACKUP) creates sequential disk save sets, it stores the save-set file in the MFD.

ANALYZE/DISK_STRUCTURE verifies all files contained in the directory structure by making comparisons to INDEXF.SYS. Any file found in INDEXF.SYS that is not traceable through the directory structure is “lost.” ANALYZE/DISK_STRUCTURE places lost files in the top-level directory SYSLOST.DIR if you specified /REPAIR in the command.

A.3.5 Core Image File, CORIMG.SYS

The core image file is not used by the operating system.

A.3.6 Volume Set List File, VOLSET.SYS

The volume set list file is used only on relative volume 1 of a volume set. The file contains a list of the labels of all the volumes in the set and the name of the volume set.

ANALYZE/DISK_STRUCTURE uses VOLSET.SYS to locate each volume in the set and confirm the attributes of each volume. Since all volume set information is stored in VOLSET.SYS on relative volume 1, ANALYZE/DISK_STRUCTURE ignores VOLSET.SYS on all other volumes.

A.3.7 Continuation File, CONTIN.SYS

The continuation file is used as the extension file identifier when a file crosses from one volume to another volume of a loosely coupled volume set. This file is used for all but the first volume of a sequential disk save set.

A.3.8 Backup Log File, BACKUP.SYS

The backup log file is reserved for future use.

A.3.9 Pending Bad Block Log File, BADLOG.SYS

The pending bad block log file contains a list of suspected bad blocks on the volume that are not listed in the bad block file.

A.3.10 Quota File, QUOTA.SYS

The quota file is a reserved file that is used by the file system to keep track of the disk usage of each UIC on a volume. If you enable disk quota checking for a volume, the records of the file QUOTA.SYS contain all the UICs on the volume. The system constantly updates QUOTA.SYS to reflect the current disk usage, the maximum allowed disk usage, and the permitted overdraft for each UIC.

During the course of its operations, ANALYZE/DISK_STRUCTURE creates a version of QUOTA.SYS in memory that reflects the actual disk usage for each UIC. This version is eventually compared to the disk version of QUOTA.SYS. If ANALYZE/DISK_STRUCTURE detects any disparities in disk usage, ANALYZE/DISK_STRUCTURE notifies you. If you invoked ANALYZE/DISK_STRUCTURE with the /REPAIR qualifier, the disk version of QUOTA.SYS is updated.

A.3.11 Volume Security Profile, SECURITY.SYS

The volume security profile includes the volume owner UIC, the volume system-owner-group-world (SOGW) protection mask, and the volume access control list (ACL).

Files–11 Disk Structure

A.4 Files–11 ODS Level 1 (VAX Only) Versus Levels 2 and 5

A.4 Files–11 ODS Level 1 (VAX Only) Versus Levels 2 and 5

On VAX systems, for reasons of performance, reliability, and security, Files–11 ODS Level 2, a compatible superset of ODS Level 1, is the preferred disk structure on the system. At volume initialization time, Structure Level 2 is the default. (Refer to the INITIALIZE command in the *OpenVMS DCL Dictionary*.)

On VAX systems, specify ODS Level 1 only for volumes that must be transportable to RSX–11M, RSX–11D, RSX–11M–PLUS, and IAS systems, as these systems support only that structure level. Additionally, you might be required to handle Structure Level 1 volumes transported to OpenVMS from one of these systems.

Where Structure Level 1 volumes are in use on the system, bear in mind the limitations on them that are shown in Table A–4.

Table A–4 Limitations on Files–11 Structure Level 1 Volumes

Disk	Only Files–11 ODS–2 disks are protected objects.
Directories	No hierarchies of directories and subdirectories, and no ordering of directory entries (that is, the file names) in any way. RSX–11M, RSX–11D, RSX–11M–PLUS, and IAS systems do not support subdirectories and alphabetical directory entries.
Disk quotas	Not supported.
Multivolume files and volume sets	Not supported.
Placement control	Not supported.
Caches	No caching of file header blocks, file identification slots, or extent entries.
System disk	Cannot be a Structure Level 1 volume.
OpenVMS Cluster access	Local access only; cannot be shared across a cluster.
Clustered allocation	Not supported.
Backup home block	Not supported.
Protection code E	E means “extend” for the RSX–11M operating system but is ignored by OpenVMS.
File versions	Limited to 32,767; version limits are not supported.
Enhanced protection features (for example, access control lists)	Not supported.
Long file names	Not supported.
RMS journaling	Not supported.
RMS execution statistics monitoring	Not supported.

Future enhancements to OpenVMS software will be based primarily on Structure Level 5; therefore, Structure Level 1 volumes might be further restricted in the future.

Tables of Time Differential Factors

The tables in this appendix show the time differential factors (TDFs) of various locations in the world. Each table contains a list of locations in a specific region. The information in the tables is believed to be accurate at the time of publication.

Note

Time zone rules are under control of each country, and are subject to change for political and other reasons. For up-to-date information, see the following web locations:

<http://swissinfo.net/cgi/worldtime/>
<http://times.clari.net.au/index.htm>

Table B-1 lists the time differential factors for Europe.

Table B-1 TDFs for Europe

Region	Standard Time TDF	Daylight Saving Time TDF
Great Britain, Ireland	0:00	+1:00
Western European Time	0:00	+1:00
Iceland	0:00	—
Middle European Time	+1:00	+2:00
Poland	+2:00	+3:00
Eastern European Time	+2:00	+3:00
Turkey	+2:00	+3:00

Table B-2 lists the time differential factors for North America.

Table B-2 TDFs for North America

Region	Standard Time TDF	Daylight Saving Time TDF
U.S./Eastern Time	-5:00	-4:00
U.S./Central Time	-6:00	-5:00
U.S./Mountain Time	-7:00	-6:00
U.S./Pacific Time	-8:00	-7:00

(continued on next page)

Tables of Time Differential Factors

Table B–2 (Cont.) TDFs for North America

Region	Standard Time TDF	Daylight Saving Time TDF
U.S./Indiana (East)	-5:00	—
U.S./Alaska	-9:00	-8:00
U.S./Arizona	-7:00	—
U.S./Navajo	-7:00	-6:00
U.S./Michigan	-5:00	-4:00
U.S./Aleutian Islands	-10:00	-9:00
U.S./Hawaii	-10:00	—
U.S./Samoa	-11:00	—
Canada/Newfoundland	-3:30	-2:30
Canada/Atlantic	-4:00	-3:00
Canada/Eastern	-5:00	-4:00
Canada/Central	-6:00	-5:00
Canada/East–Saskatchewan	-6:00	—
Canada/Mountain	-7:00	-6:00
Canada/Pacific	-8:00	-7:00
Canada/Yukon	-9:00	-8:00

Table B–3 lists the time differential factors for Central and South America.

Table B–3 TDFs for Central and South America

Region	Standard Time TDF	Daylight Saving Time TDF
Mexico/BajaNorte	-8:00	-7:00
Mexico/BajaSur	-7:00	—
Mexico/General	-6:00	—
Cuba	-5:00	-4:00
Jamaica	-5:00	-4:00
Brazil/East	-3:00	-2:00
Brazil/West	-4:00	-3:00
Brazil/Acre	-5:00	-4:00
Brazil/DeNoronha	-2:00	-1:00
Chile/Regional	-4:00	-3:00
Chile/Easter Island	-6:00	-5:00

Tables of Time Differential Factors

Table B–4 lists the time differential factors for Asia.

Table B–4 TDFs for Asia

Region	Standard Time TDF	Daylight Saving Time TDF
PRC (Mainland China)	+8:00	+9:00
ROK (Korea)	+9:00	+10:00
Israel	+2:00	+3:00
Iran	+3:30	+4:30
Japan	+9:00	—
Singapore	+8:00	—
Hong Kong	+8:00	—
ROC (Taiwan)	+8:00	—

Table B–5 lists the time differential factors for the South Pacific.

Table B–5 TDFs for the South Pacific

Region	Standard Time TDF	Daylight Saving Time TDF
Australia/Tasmania	+10:00	+11:00
Australia/Queensland (standard time only)	+10:00	—
Australia/Queensland	+10:00	+11:00
Australia/North	+9:30	—
Australia/West	+8:00	—
Australia/South	+9:30	+10:30
Australia/Victoria	+10:00	+11:00
Australia/New South Wales	+10:00	+11:00
New Zealand	+12:00	+13:00

Table B–6 lists the time differential factors for Antarctica.

Table B–6 TDFs for Antarctica

Region	Standard Time TDF	Daylight Saving Time TDF
Antarctica	+0:00	—

Compaq MIB Subagents Implemented on OpenVMS Alpha

The Extensible Simple Network Management Protocol (eSNMP) allows network managers to manage many different types of devices across all network and vendor boundaries through the use of databases called Management Information Bases (MIBs). Essentially, information is exchanged between **master agents** and **subagents**, which are devices such as routers and servers on the network being managed, and **managers**, which are the devices on the network through which the management is done.

This appendix describes the Compaq Server MIB and the Compaq Cluster MIB.

C.1 Compaq Server MIB Subagents

The Compaq Server MIB (DSM) consists of two extensions, or subagents:

Extension	Describes
System	A management interface to Alpha system information not defined by standard MIBs
Management	Instrumentation in the Compaq extension agent, including the ability to detect and monitor thresholds on integer variables

The representation of the DSM within the standard Structure of Managed Information (SMI) framework is:

```
iso(1) org(3) dod(6) internet(1) private(4) enterprises(1) 36
```

OpenVMS Alpha Version 7.1-1H1 and later implements the DSM subagents on the AlphaServer 800, 1000, 4000, 4100, 8200, and 8400 systems. With the DSM subagents, customers can remotely determine and manage important information such as:

- Firmware revision numbers
- Base system descriptions
- Field Replaceable Unit (FRU) information and descriptions
- Processor and cache status
- Interface configurations
- Environmental conditions in the system enclosure that might be detrimental to the hardware

Use the following software to access the DSM subagents:

- The DIGITAL ServerWORKS Manager Version 3.0 or any MIB browser that has access to the DSM definitions.

Compaq MIB Subagents Implemented on OpenVMS Alpha

C.1 Compaq Server MIB Subagents

- DIGITAL TCP/IP Services for OpenVMS Version 4.1 or later. The DSM subagents use the SNMP agent supplied with TCP/IP Services to communicate with SNMP clients.

The following sections describe the DSM subagents and explain how to set up your system to use them.

C.1.1 Overview of DSM Subagents

DSM subagents respond to SNMP requests for a DSM **object** — the data item that the network manager is concerned with, or a **trap** — information about a change of status. A subagent is responsible for reporting on and maintaining the data pertaining to these objects and traps.

The DSM system subagent implements the objects listed in Table C–1. Each object corresponds to a group of base system and environmental information relevant to OpenVMS Alpha networking and can be accessed by a network manager through ServerWORKS Manager.

Table C–1 DSM System Subagent Objects Implemented on OpenVMS Alpha

Object	Data Type	Access	Description
MIB Information Group			
svrSysMibMajorRev	Integer	Read only	The major revision number of this implementation of the svrSystem MIB. Currently 1.
svrSysMibMinorRev	Integer	Read only	The minor revision number of this implementation of the svrSystem MIB. Currently 0.
Base System Description Group			
svrSystemModel	DisplayString	Read only	System and model name. For example, AlphaServer 2100.
svrSystemDescr	DisplayString	Read only	General text description of system type.
svrSystemBoardFruIndex	Integer	Read only	The index of the Field Replaceable Unit (FRU) in the FRU table describing the serial number and other asset information of the board. If unknown, 0.
svrSystemBootedOS	Integer	Read only	The current booted operating system.
svrSystemShutdownReason	DisplayString	Read only	The possible reason for the system shutdown.
svrFirmwareIndex	Integer	Read only	An index value unique to the local system.
svrFirmwareDescr	DisplayString	Read only	Descriptive text for items such as the SRM console, ARC console, and system BIOS.
svrFirmwareRev	DisplayString	Read only	A version number, often of the form Vx.y or Vx.y-z.
svrFwSymbolName	DisplayString	Read only	The symbol name as visible on the console.
svrFwSymbolValue	Octet string	Read only	The symbol value. Null if none or unknown.

(continued on next page)

Compaq MIB Subagents Implemented on OpenVMS Alpha C.1 Compaq Server MIB Subagents

Table C–1 (Cont.) DSM System Subagent Objects Implemented on OpenVMS Alpha

Object	Data Type	Access	Description
System Processor Group			
svrCpuIndex	Integer	Read only	An index value for the CPU entry that is unique to the local system.
svrCpuManufacturer	DisplayString	Read only	The manufacturer of the processor.
svrCpuRevision	DisplayString	Read only	Version information in processor-specific format.
svrCpuFruIndex	Integer	Read only	The index of the FRU entry in the FRU table that describes the asset information of the component containing the processor. If unknown, 0.
svrCpuCacheIndex	Integer	Read only	The local index value.
svrCpuCacheLevel	Integer	Read only	Level 1, level 2, level 3 cache, other, or unknown.
svrCpuCacheType	Integer	Read only	Type of cache: internal, external, internal instruction, or internal data.
svrCpuCacheSize	Kbytes	Read only	Cache size in Kbytes.
svrCpuCacheSpeed	Integer	Read only	Cache speed in nanoseconds. If unknown, 0.
svrCpuCacheStatus	Integer	Read only	Current status of the cache: enabled, disabled, other, or unknown.
Memory Configuration Group			
svrPhysicalMemorySize	Kbytes	Read only	Total amount of physical memory as seen by the operating system.
svrPhysicalMemoryFree	Kbytes	Read only	Amount of free physical memory.
svrMemIndex	Integer	Read only	Unique index for this entry.
svrMemSize	Kbytes	Read only	Length of memory range.
svrMemFruIndex	Integer	Read only	Index of the FRU entry in the FRU table on which the memory resides. If unknown, 0.
svrBusIndex	Integer	Read only	An index value that is unique to the local system.
svrBusType	BusTypes	Read only	Bus type.
svrLogicalSlotNumber	Integer	Read only	Unique logical slot number on a given bus.
svrLogicalSlotDescr	DisplayString	Read only	Device description derived from ID or as set by the management station.
svrLogicalSlotRevision	DisplayString	Read only	Vendor-supplied major and minor revision of device in the slot.

(continued on next page)

Compaq MIB Subagents Implemented on OpenVMS Alpha

C.1 Compaq Server MIB Subagents

Table C–1 (Cont.) DSM System Subagent Objects Implemented on OpenVMS Alpha

Object	Data Type	Access	Description
Physical Configuration Group			
svrFruIndex	Integer	Read only	An index value that is unique to the system.
svrFruType	Integer	Read only	General category of FRU type.
svrFruDescr	DisplayString	Read only	Detailed description of the FRU type, if known.
svrFruVendor	DisplayString	Read only	Manufacturer's name or ID.
svrFruPartNumber	DisplayString	Read only	Order number for this unit.
svrFruRevision	DisplayString	Read only	Version number of the unit. If an illustration is available, it appears as "Artwork: XXX" following the FRU version number.
svrFruFirmwareRevision	DisplayString	Read only	Revision of the firmware, if applicable. Otherwise, null.
svrFruSerialNumber	DisplayString	Read only	Unit's serial number.
Environment Group: Thermal			
svrThermalSensorCount	Integer	Read only	Number of thermal sensors present and readable in the system.
svrThSensorIndex	Integer	Read only	An index value unique to the local system.
svrThSensorReading	Integer	Read only	Current value read by the sensor in units as described by the svrThSensorReadingUnits object.
svrThSensorReadingUnits	ThermUnits	Read only	Value of sensor in degrees Fahrenheit, Celsius, or relative value. If not available, value will be unknown.
svrThSensorStatus	Integer	Read only	The sensor's status value.
Environment Group: Cooling			
svrFanCount	Integer	Read only	The number of fans whose states are detectable.
svrFanIndex	Integer	Read only	An index value unique to the local system.
svrFanStatus	Integer	Read only	Current fan status.
Environment Group: Power Supply			
svrPowerSupplyCount	Integer	Read only	Number of detectable power supplies reflected as entries in the svrPowerSupplyTable object.
svrPowerSupplyIndex	Integer	Read only	An index value unique to the local system.
svrPowerSupplyStatus	Integer	Read only	Current state of the power supply.

The DSM Management subagent implements the objects and traps listed in Tables C–2 and C–3, respectively.

Compaq MIB Subagents Implemented on OpenVMS Alpha

C.1 Compaq Server MIB Subagents

Each object or trap corresponds to a group of management areas relevant to OpenVMS Alpha networking and can be accessed by a network manager through ServerWORKS Manager.

Table C-2 DSM Management Subagent Objects Implemented on OpenVMS Alpha

Object	Data Type	Access	Description
MIB Information Group			
svrMgtMibMajorRev	Integer	Read only	The major revision number of this implementation of the svrMgt MIB. Currently 1.
svrMgtMibMinorRev	Integer	Read only	The minor revision number of this implementation of the svrMgt MIB. Currently 0.
Alarms Group			
svrAlarmNextThrIndex	Integer	Read only	The next available index for creating a svrThrEntry object. If the value is -1, the maximum number of thresholds has been reached. A threshold record cannot be created until you delete the current threshold record.
svrAlarmEnableTraps	Boolean	Read/write	If true, a trap is sent for each triggered alarm.
svrThresholdTable	Sequence of SvrThresholdEntry	Not accessible	<p>The threshold table that describes conditions for setting and resetting alarms. The agent checks this table for exceptions.</p> <p>You can set alarms on absolute values (such as the current integer value of the sampled variable) or on delta values (such as the difference between the current or last value). Alarms can be <i>Greater Than</i> exception alarms, <i>Less Than</i> exception alarms, <i>Equal To</i> alarms, and so on. (See the svrThrAlarmType object description.)</p> <p>Hysteresis (the tendency of certain binary devices to show different threshold values when changing from 0 to 1 than when changing from 1 to 0) is introduced by providing thresholds both for setting and resetting of the alarm state, thereby limiting the number of traps that are sent on alarm triggering.</p> <p>You can create alarms to persist across agent reboots; however, this is not recommended for dynamic table variables.</p> <p>The triggering of an alarm changes a state variable in the conceptual row and can also trigger the sending of a trap, or the local logging of an event.</p>

(continued on next page)

Compaq MIB Subagents Implemented on OpenVMS Alpha

C.1 Compaq Server MIB Subagents

Table C–2 (Cont.) DSM Management Subagent Objects Implemented on OpenVMS Alpha

Object	Data Type	Access	Description
Alarms Group			
svrThresholdEntry	SvrThresholdEntry	Not accessible	<p>A threshold alarm set on an integer variable.</p> <p>An alarm entry is created by the management console using the current value of svrAlarmNextThrIndex to name the instances of the row variables, setting the svrThrStatus to underCreation. When you create a threshold entry for the first time, issue a set request on svrThrStatus.</p> <p>You can set the remaining row variables in the same operation or in subsequent operations. Those not set retain their default values as described. You must set variable values for the following objects in the Alarms group before you enable the alarm:</p> <ul style="list-style-type: none"> • svrThrStatus (set to underCreation) • svrThrVariableName through svrThrSeverity (set appropriately; see the object descriptions)
svrThrIndex	Integer	Read only	<p>An index value unique to the local system. On creation, set to the value of svrAlarmNextThrIndex.</p>
svrThrStatus	Integer	Read/write	<p>Describes the status of the row.</p> <p>When the row is created with the initial set, you must set svrThrStatus to underCreation. When the management console has completed the row setup, it sets this variable to rowEnabled. Variables in the row can only be written if svrThrStatus is in the initial underCreation state or has been set to rowDisabled.</p> <p>To delete the row, set the status to rowInvalid. Be aware that errors in variable polling and threshold checking that cannot be corrected cause a row status change to rowError. Once the status is set to rowError by the agent, the agent does not reset the status. Instead, the management console must reset the status based on information returned with svrThrErrorValue or for other reasons.</p>
svrThrVariableName	Object identifier	Read/write	<p>The object identifier (OID) of an integer variable to be tested against the threshold. At row creation, the variable equals the value 0.0 and must be set to the OID of an integer variable before enabling the alarm.</p>

(continued on next page)

Compaq MIB Subagents Implemented on OpenVMS Alpha C.1 Compaq Server MIB Subagents

Table C-2 (Cont.) DSM Management Subagent Objects Implemented on OpenVMS Alpha

Object	Data Type	Access	Description
Alarms Group			
svrThrValueType	Integer	Read/write	Absolute or delta value. The default on row creation is absoluteValue. The deltaValue is calculated by taking the current value and subtracting the svrThrLastValue value.
svrThrAlarmType	Integer	Read/write	<p>An alarm that signals a threshold whose value is <i>Greater Than</i>, <i>Greater Than or Equal To</i>, <i>Equal To</i>, <i>Less Than or Equal To</i>, or <i>Less Than</i>. The default value on row creation is <i>Greater Than</i>.</p> <p><i>Greater Than</i> or <i>Greater Than or Equal To</i> thresholds for absolute values occur when the sample value equals or exceeds the svrThrThresholdValue and svrThrAlarmState was reset. This condition causes svrThrAlarmState to be set and, if svrAlarmEnableTraps is true, a svrThrExceptTrap is sent. SvrThrAlarmState is reset when the sample value falls below or equals svrThrResetValue.</p> <p><i>Less Than</i> or <i>Less Than or Equal To</i> thresholds for absolute values occur when the sample value falls below or equals the svrThrThresholdValue, and svrThrAlarmState was reset. This condition causes the svrThrAlarmState to be set and, if svrAlarmEnableTraps is true, a svrThrExceptTrap is sent. SvrThrAlarmState is reset when the sample value exceeds or equals svrThrResetValue.</p> <p><i>Equal To</i> thresholds for absolute values occur when the sample value equals svrThrThresholdValue and svrThrAlarmState was reset. This condition causes the svrThrAlarmState to be set and, if svrAlarmEnableTraps is true, a svrThrExceptTrap is sent. SvrThrAlarmState is reset when the sample value does not equal svrThrResetValue.</p> <p>The same conditions apply for delta values as for absolute values except the difference between the sample value and the svrThrLastValue is used for comparison with both the svrThrThresholdValue and the svrThrResetValue. Note that it is possible to have negative delta values since the difference is computed as the current value minus the svrThrLastValue.</p>

(continued on next page)

Compaq MIB Subagents Implemented on OpenVMS Alpha

C.1 Compaq Server MIB Subagents

Table C-2 (Cont.) DSM Management Subagent Objects Implemented on OpenVMS Alpha

Object	Data Type	Access	Description
Alarms Group			
svrThrSampleInterval	Integer	Read/write	The interval (in seconds) between polls to check for threshold exceptions. The default value on row creation is 30 seconds. Minimum value: 1.
svrThrPersistent	Boolean	Read/write	If true, the threshold persists across agent restarts. Default on row creation: false. By default, the files used to store persistent data are SYS\$SYSTEM:TCPIP\$MGT_THRESHOLDS.DAT and SYS\$SYSTEM:TCPIP\$MGT_THRESHOLDS.BAK. To move the files off the system disk or rename them, the system manager can define the logical names TCPIP\$MGT_PERSISTENCE_DAT and TCPIP\$MGT_PERSISTENCE_BAK in the SYS\$MANAGER:SYLOGICALS.COM file as appropriate. For example, to point to files in a different location, add the following definitions to SYS\$MANAGER:SYLOGICALS.COM. (The examples are formatted to fit the table column.) <pre> \$ DEFINE/SYS - _\$ TCPIP\$MGT_PERSISTENCE_DAT - _\$ DISK2:[SNMP.MIB]PERSIST.DAT; \$ DEFINE/SYS - _\$ TCPIP\$MGT_PERSISTENCE_BAK - _\$ DISK2:[SNMP.MIB]PERSIST.BAK; </pre>
svrThrThresholdValue	Integer	Read/write	The threshold value that is compared to the current or delta value. Default on row creation: 0.
svrThrResetValue	Integer	Read/write	The value used to reset the threshold on all svrThrAlarmTypes objects except those that are <i>Equal To</i> . Default on row creation: 0.
svrThrLastValue	Integer	Read only	The previous sample needed to evaluate if alarm should be triggered or to evaluate delta values for threshold checking.
svrThrAlarmState	Integer	Read only	Indicates whether the alarm is currently set or reset. Used by polling management applications to determine if a threshold exception state has been detected based on the alarm definition. Has an initial value of reset when the alarm is enabled or the agent is restarted. The value is reset if svrThrStatus changes to rowDisabled or rowInvalid. For guidelines on state changes, see the description for svrThrAlarmType.

(continued on next page)

Compaq MIB Subagents Implemented on OpenVMS Alpha C.1 Compaq Server MIB Subagents

Table C–2 (Cont.) DSM Management Subagent Objects Implemented on OpenVMS Alpha

Object	Data Type	Access	Description
Alarms Group			
svrThrLogEvent	Boolean	Read/write	If true, logs data to the subagent process log file; for example, to [TCPIP\$SNMP]TCPIP\$SVRMGT_MIB.LOG. (See Section C.1.2.) Default value: false.
svrThrDescr	DisplayString	Read/write	Describes the type of threshold. Set by the management console, not by the agent.
svrThrErrorValue	SnmpErrors	Read only	The SNMP-defined error status that caused the svrThrStatus value to become equal to rowError. Valid only at that time.
svrThrComparisonName	Object identifier	Read/write	An object identifier (OID) to a descriptor attribute used with the svrThrPersistent value to verify that the svrThrVariableName instance is correct. Optional. Default: 0.0. On agent restarts, the value is retrieved and compared to the svrThrComparisonValue. If not equal, the OID instancing for svrThrVariableName might be incorrect. If this situation occurs, svrThrStatus is set to rowError and svrThrErrorValue to badValue.
svrThrComparisonValue	DisplayString	Read/write	Date value of svrThrComparisonName. Optional. Used when svrThrPersistent is set. The value is compared to the current value on agent restarts. Default: null.
svrThrSeverity	Severity	Read/write	Indicates the severity of the threshold. Default on row creation: informational.

Table C–3 DSM Management Subagent Traps Implemented on OpenVMS Alpha

Trap	Variable	Description
Local Server Control Group		
svrThrHighExceptTrap	svrThrVariableName svrThrValueType svrThrThresholdValue svrThrLastValue svrThrDescr	A high severity trap. The value that caused the alarm to occur is returned in svrThrLastValue.
svrThrMediumExceptTrap	svrThrVariableName svrThrValueType svrThrThresholdValue svrThrLastValue svrThrDescr	A medium severity trap. The value that caused the alarm to occur is returned in svrThrLastValue.

(continued on next page)

Compaq MIB Subagents Implemented on OpenVMS Alpha

C.1 Compaq Server MIB Subagents

Table C-3 (Cont.) DSM Management Subagent Traps Implemented on OpenVMS Alpha

Trap	Variable	Description
Local Server Control Group		
svrThrLowExceptTrap	svrThrVariableName svrThrValueType svrThrThresholdValue svrThrLastValue svrThrDescr	A low severity trap. The value that caused the alarm to occur is returned in svrThrLastValue.
svrThrInformationalExceptTrap	svrThrVariableName svrThrValueType svrThrThresholdValue svrThrLastValue svrThrDescr	An Informational severity trap. The value that caused the alarm to occur is returned in svrThrLastValue.

C.1.2 Setting Up the System to Use the DSM Subagents

To configure SNMP on the system and enable the master agent to accept SET commands from SNMP clients, issue the following TCP/IP Services management command from the TCPIP> prompt. This operation requires SYSPRV or BYPASS privileges.

```
TCPIP> SET CONFIGURATION SNMP /FLAGS=SETS
```

To enable or disable the type of access to your local MIB data, use the following commands, qualifiers, and options:

```
TCPIP> SET CONFIGURATION SNMP /[NO]COMMUNITY="name" -
_TCPIP> /[NO]ADDRESS=host address -
_TCPIP> /TYPE=[NO]READ,[NO]TRAP,[NO]WRITE
```

For example, the following command configures SNMP, specifies the community name and address, specifies that the agent can accept SET commands from members of the community, and enables the master agent to send trap messages to members of the community. (Note that READ access is assumed when specifying TRAP or WRITE.)

```
TCPIP> SET CONFIGURATION SNMP /COMMUNITY="public" /ADDRESS=128.45.2.8 -
_TCPIP> /TYPE=TRAP,WRITE
```

For your convenience, the following files have example entries to start, run, and shut down the DSM subagents.

- [TCPIP\$SNMP]TCPIP\$EXTENSION_MIB_STARTUP.COM. (This file is called by TCPIP\$SNMP_STARTUP.COM.)
- [TCPIP\$SNMP]TCPIP\$EXTENSION_MIB_RUN.COM. (This file is called by TCPIP\$SNMP_RUN.COM.)
- [TCPIP\$SNMP]TCPIP\$EXTENSION_MIB_SHUTDOWN.COM. (This file is called by TCPIP\$SNMP_SHUTDOWN.COM.)

Search the files for SVRSYSTEM_MIB entries and edit them as needed. You must also remove the GOTO entries, which cause the command procedures to simply exit.

C.2 Compaq Cluster MIB Subagents

The Compaq Cluster MIB (DCM) is a private Compaq management information base that delivers management information about OpenVMS Cluster systems.

The DCM consists of two extensions, or subagents:

Extension	Describes
System	A management interface to cluster system information not defined by standard MIBs
Management	Instrumentation in the Compaq extension agent, including the ability to detect and monitor thresholds on integer variables

The representation of the DCM within the standard Structure of Managed Information (SMI) framework is:

```
iso(1) org(3) dod(6) internet(1) private(4) enterprises(1) 36
```

OpenVMS Alpha Version 7.2 implements the DCM subagents. With the DCM subagents, you can remotely determine status information for an OpenVMS Cluster system, including the following:

- Cluster software version
- Status of the cluster software: installed, running, failed, and so forth
- Type of cluster that is running
- State change to reflect when a cluster member is added or deleted

Use the following software to access the DCM subagents:

- The DIGITAL ServerWORKS Manager Version 3.0 or any MIB browser that has access to the DCM definitions.
- DIGITAL TCP/IP Services for OpenVMS Version 4.1 or later. The DCM subagents use the SNMP agent supplied with TCP/IP Services to communicate with SNMP clients.

The following sections describe the DCM subagents and explain how to set up your system to use them.

C.2.1 Overview of DCM Subagents

DCM subagents respond to SNMP requests for a DCM **object** — the data item that the network manager is concerned with, or a **trap** — information about a change of status. A subagent is responsible for reporting on and maintaining the data pertaining to these objects and traps.

The DCM subagents implement the objects listed in Table C-4. Each object returns information relevant to an OpenVMS Cluster system and can be accessed by a network manager through ServerWORKS Manager.

Compaq MIB Subagents Implemented on OpenVMS Alpha

C.2 Compaq Cluster MIB Subagents

Table C-4 DCM Subagent Objects Implemented on OpenVMS

Object	Data Type	Access	Description
Cluster Information			
svrCluSoftwareVendor	DisplayString	Read only	Cluster software vendor name. The current value is Digital.
svrCluSoftwareVersion	DisplayString	Read only	Cluster software version. This is the OpenVMS version string.
svrCluSoftwareStatus	ClusterStatus	Read only	The status of the cluster software. Possible values are running and not running.
svrCluClusterType	ClusterType	Read only	The type of cluster that is running. The current value is OpenVMS.
svrCluExtensionOID	Object Identifier	Read only	The authoritative identification for the MIB. If no such identifier exists, the value {0.0} is returned.
svrCluThisMember	Integer	Read only	Index into the member table (svrCluMemberTable) that corresponds to this node.
SNMP Traps			
svrCluMemberAdded	Trap Packet	Read only	Generated when a cluster member is added.
svrCluMemberDeleted	Trap Packet	Read only	Generated when a cluster member is deleted.

(continued on next page)

Compaq MIB Subagents Implemented on OpenVMS Alpha C.2 Compaq Cluster MIB Subagents

Table C–4 (Cont.) DCM Subagent Objects Implemented on OpenVMS

Object	Data Type	Access	Description
Node-Specific Information			
svrCluMemberIndex	Integer	Read only	A unique index for the entry. Values of svrCluMemberIndex must remain constant at least between reboots of the network management system on the managed node. Where possible, this value should reflect the system's native member identifier.
svrCluMemberName	DisplayString	Read only	The SCS node name of this cluster member. A zero-length value means the member's node name is unknown. This name may not necessarily resolve to an address.
svrCluMemberComment	DisplayString	Read only	This is the hardware name of the node, as returned by the \$GETSYI system service.
svrCluMemberStatus	MemberStatus	Read only	Status of this member. Possible values are normal and removed.
svrCluMemberAddressIndex	Integer	Read only	The index for this address.
svrCluMemberAddress	IpAddress	Read only	An IP address of this cluster member. This address may not be reachable from nodes that are not configured into the cluster.

C.2.2 Setting Up the System to Use the DCM Subagents

For your convenience, the following files have example entries to start, run, and shut down the DCM subagents.

- [TCPIP\$SNMP]TCPIP\$EXTENSION_MIB_STARTUP.COM. (This file is called by TCPIP\$SNMP_STARTUP.COM.)
- [TCPIP\$SNMP]TCPIP\$EXTENSION_MIB_RUN.COM. (This file is called by TCPIP\$SNMP_RUN.COM.)
- [TCPIP\$SNMP]TCPIP\$EXTENSION_MIB_SHUTDOWN.COM. (This file is called by TCPIP\$SNMP_SHUTDOWN.COM.)

Search the files for SVRCLUSTER_MIB entries and edit them as needed. You must also remove the GOTO entries, which cause the command procedures to simply exit.

Managing the OpenVMS Registry

The OpenVMS Registry Server manages the OpenVMS Registry Database.

For the latest information about the OpenVMS Registry, see the *OpenVMS Connectivity Developer Guide*. This document is available as part of the *COM for OpenVMS* kit and from the OpenVMS website (www.compaq.com/openvms) in the OpenVMS documentation area.

This chapter describes how to install and manage the OpenVMS Registry server, which is part of the OpenVMS Version 7.3 system installation.

D.1 Accessing the OpenVMS Registry

The OpenVMS Registry database is created the first time you start the OpenVMS Registry server. You must define the `SYS$REGISTRY` logical before starting the OpenVMS Registry server. The OpenVMS Registry server will not start if you have not defined the `SYS$REGISTRY` logical.

To access the OpenVMS Registry using the COM APIs, you must install COM for OpenVMS.

To access the OpenVMS Registry using the Windows NT application `RegEdt32`, you must first install, configure, and start Advanced Server for OpenVMS. For more information, see the Advanced Server for OpenVMS documentation.

You can also access the OpenVMS Registry using the OpenVMS Registry server management utility or the OpenVMS Registry system services, which are installed as part of the OpenVMS Registry in OpenVMS Version 7.2-1 and higher.

D.2 Configuring the OpenVMS Registry: the `REG$CONFIG` Configuration Utility

The OpenVMS Registry Configuration utility (`REG$CONFIG`) provides information about the OpenVMS Registry server status and the OpenVMS Registry database location, and allows you to change OpenVMS Registry logical names and paths.

Enter the following command to start the OpenVMS Registry Configuration utility:

```
$ @SYS$MANAGER:REG$CONFIG
```

The system displays the following menu:

Managing the OpenVMS Registry

D.2 Configuring the OpenVMS Registry: the REG\$CONFIG Configuration Utility

```
-----  
OpenVMS Registry Configuration Utility  
-----  
1 - Configure OpenVMS Registry logical names and directory paths  
2 - Display OpenVMS Registry logical names and directory paths  
3 - Check the state of the OpenVMS Registry server  
4 - Start the OpenVMS Registry server on this node  
H - Help about this utility  
[E] - Exit  
Please enter your choice :  
-----
```

To select an option, enter the option number. The options are as follows:

- 1 - Configure OpenVMS Registry logical names and directory paths
Allows you to configure the OpenVMS Registry server startup value and specify the location of the OpenVMS Registry database.
For this procedure, see Section D.2.1.
- 2 - Display OpenVMS Registry logical names and directory paths
Displays the current values of the OpenVMS Registry server logical (startup value) for this node and the OpenVMS Registry database location.
- 3 - Check the state of the OpenVMS Registry server
Displays the current state of the OpenVMS Registry server. The system displays one of the following:
 - The OpenVMS Registry server is started in the cluster.
 - The OpenVMS Registry server is started on this node.
 - The OpenVMS Registry server is not started.
- 4 - Start the OpenVMS Registry server on this node
Starts the OpenVMS Registry server on the current node. The system displays the following message:
 - The OpenVMS Registry server has successfully started.
- H - Help about this utility
Displays online help for OpenVMS Registry Configuration utility options.
- [E] - Exit
Exits the OpenVMS Registry Configuration utility.

----- **Tip: Enter Q (Quit) at any time** -----

You can enter Q at any prompt to return to the OpenVMS Registry Configuration utility menu.

If you quit while you are configuring logical names, the system updates only those values for which you have received a confirmation message.

D.2 Configuring the OpenVMS Registry: the REG\$CONFIG Configuration Utility

D.2.1 Configuring OpenVMS Registry Values

The system displays the following questions:

1. The system prompts you to enter standalone or cluster information. The system displays the following message:

Is this system now a node in a cluster or will this system become part of a cluster? (Y/N/Q):

2. The system displays the current information about the REG\$TO_BE_STARTED logical, then prompts you to change the value.

- REG\$TO_BE_STARTED -

[current value of REG\$TO_BE_STARTED]

NOTE: Setting this logical to TRUE starts the OpenVMS Registry server automatically when the system boots. Setting this logical to FALSE prevents the OpenVMS Registry server from starting when the system boots and prevents other products from starting the OpenVMS Registry server. If the OpenVMS Registry Server is not started at boot time, but other products that require an OpenVMS Registry server are able to start the OpenVMS Registry server, you do not need to assign a value to this logical.

Do you want to change this value? (Y/N/Q) [Y]:

If you choose Y, the system prompts you for the new value.

Enter the new value (TRUE/FALSE/NOVAL/Q):

Enter one of the following:

Action	Value
Start the OpenVMS Registry server on reboot. Allow other products to start the server.	TRUE
Do not start the OpenVMS Registry server on reboot. Do not allow other products to start the server.	FALSE
Do not start the OpenVMS Registry server on reboot. Allow other products to start the server. (Deassigns the logical name.)	NOVAL
Quit this procedure and return to the OpenVMS Registry Configuration utility menu.	Q

In which logical name table do you want the logical defined? (SYSTEM/SYSCLUSTER/CLUSTER/Q) :

Enter one of the following:

Action	Value
Add the REG\$TO_BE_STARTED logical to the LNM\$SYSTEM logical name table. This table contains names that are shared by all processes in the system.	SYSTEM
Add the REG\$TO_BE_STARTED logical to the LNM\$SYSCLUSTER logical name table. This table contains names that are shared by all processes in an OpenVMS Cluster.	SYSCLUSTER
Add the REG\$TO_BE_STARTED logical to the LNM\$CLUSTER logical name table. This table is the parent table for all clusterwide logical name tables.	CLUSTER
Quit this procedure and return to the OpenVMS Registry Configuration utility menu.	Q

Managing the OpenVMS Registry

D.2 Configuring the OpenVMS Registry: the REG\$CONFIG Configuration Utility

After you enter the new or updated value, the system confirms the change and displays the line you must add to your SYLOGICALS.COM file.

```
The logical REG$TO_BE_STARTED has been temporarily defined.
Before you reboot the system you must edit your SYLOGICALS.COM
to include the line:
```

```
DEFINE/TABLE=table-name REG$TO_BE_STARTED value
```

Press [Enter] to continue.

3. The system displays the current information about the SYS\$REGISTRY logical, then prompts you to change the value.

```
- SYS$REGISTRY logical -
```

```
current value of SYS$REGISTRY
```

```
Note: When the OpenVMS Registry server is started, the system
creates an OpenVMS Registry database at this location.
If an OpenVMS Registry database already exists on your system,
you must redefine the SYS$REGISTRY logical to point to the
existing OpenVMS Registry database location.
```

Do you wish to change this value? (Y/N/Q) [Y]:

If you choose Y, the system prompts you for the new value.

Enter the new value for SYS\$REGISTRY ("yourvalue"/NOVAL/Q):

Enter one of the following:

Action	Value
Define a new or changed location for the OpenVMS Registry database.	A valid directory specification, such as DKA0: [SYS\$REGISTRY]
Deassign the logical name.	NOVAL
Quit this procedure and return to the OpenVMS Registry Configuration utility menu.	Q

4. The system displays your updated value and prompts you to confirm the value.

```
You have entered: value
Is this correct? (Y/N/Q) [Y]:
```

5. The system prompts you to enter a logical table name in which to store the new or updated logical.

```
In which logical name table do you want the logical defined?
(SYSTEM/SYSCLUSTER/CLUSTER/Q):
```

Enter one of the following:

D.2 Configuring the OpenVMS Registry: the REG\$CONFIG Configuration Utility

Action	Value
Add the SYS\$REGISTRY logical to the LNM\$SYSTEM logical name table. This table contains names that are shared by all processes in the system.	SYSTEM
Add the SYS\$REGISTRY logical to the LNM\$SYSCLUSTER logical name table. This table contains names that are shared by all processes in an OpenVMS Cluster.	SYSCLUSTER
Add the SYS\$REGISTRY logical to the LNM\$CLUSTER logical name table. This table is the parent table for all clusterwide logical name tables.	CLUSTER
Quit this procedure and return to the OpenVMS Registry Configuration utility menu.	Q

After you enter the new or updated value, the system confirms the change and displays the line you must add to your SYLOGICALS.COM file.

The logical SYS\$REGISTRY has been temporarily defined.
Before you reboot the system you must edit your SYLOGICALS.COM file to include the line:

```
DEFINE/TABLE=table-name SYS$REGISTRY dir-spec
```

Press [Enter] to continue.

- The system displays information about the location of the OpenVMS Registry database.

```
- SYS$REGISTRY directory -
```

```
[directory status]
```

If the directory does not exist, the system prompts you to create the directory.

!!Caution!! When the OpenVMS Registry server starts, the system creates an OpenVMS Registry database at this location. If you already have an OpenVMS Registry database on your system, you must redefine the SYS\$REGISTRY logical to point to that location.

Do you wish to create the directory? (Y/N/Q) [Y]:

If you enter Y the system confirms the directory creation.

The SYS\$REGISTRY directory has now been created.

Press [Enter] to return to the menu.

D.3 Starting the OpenVMS Registry

You can control how the OpenVMS Registry will start as follows:

- Start the OpenVMS Registry automatically when the system reboots.
- Have products that require the OpenVMS Registry to be running start the OpenVMS Registry.
- Start the OpenVMS Registry manually.
- Prevent the OpenVMS Registry from starting.

Use the OpenVMS Registry Configuration utility described in Section D.2 to control how the OpenVMS Registry starts.

Managing the OpenVMS Registry

D.3 Starting the OpenVMS Registry

D.3.1 Starting the OpenVMS Registry Manually

Under some conditions, you might want to start the OpenVMS Registry server manually.

Compaq recommends that you use the `SYS$STARTUP:REG$STARTUP.COM` command procedure. The following command procedure ensures that the server process quotas are set to the required minimum values:

```
$ @SYS$STARTUP:REG$STARTUP.COM
```

Alternately, you can use the following command to start the OpenVMS Registry manually:

```
$ SET SERVER REGISTRY_SERVER/START
```

D.4 Shutting Down the OpenVMS Registry

The OpenVMS Registry server is shut down automatically as part of a system shutdown.

If you want to shut down the OpenVMS Registry manually, use the following command:

```
$ SET SERVER REGISTRY_SERVER/EXIT
```

D.5 OpenVMS Registry Server Commands

The OpenVMS Registry server commands allow you to display (SHOW) and change (SET) the state of the OpenVMS Registry server. The following sections list and describe the OpenVMS Registry server commands.

SHOW SERVER REGISTRY_SERVER

Show the current status of the OpenVMS Registry on a specified node.

This command requires the SYSPRV privilege.

Format

```
SHOW SERVER REGISTRY_SERVER  
[/MASTER | /CLUSTER | /NODE=(node,...)]  
[/PAGE]
```

Qualifiers

/MASTER

Displays the node and process ID (PID) of the current OpenVMS Registry master server in the cluster. This command does not communicate with the OpenVMS Registry servers in the cluster. Requires SYSLCK privilege as well as the SYSPRV privilege.

/CLUSTER

Returns the show output from each OpenVMS Registry server in the cluster, listing the OpenVMS Registry master server information first.

/NODE=(node,...)]

Returns OpenVMS Registry server information about the servers on the specified nodes, listed in the order in which you enter the node names. The node names you specify must be in the current cluster.

/PAGE

Displays the returned show output in a scrollable page display.

SET SERVER REGISTRY_SERVER

SET SERVER REGISTRY_SERVER

Change the state of the OpenVMS Registry.
This command requires the SYSPRV privilege.

Format

```
SET SERVER REGISTRY_SERVER  
[/MASTER | /CLUSTER | /NODE=(node,...)]  
[/START | /RESTART | /EXIT | /ABORT ]  
[/[NO]LOG ]
```

Qualifiers

/MASTER

Issues the specified command to the OpenVMS Registry master server only.
Requires the SYSLCK privilege as well as the SYSPRV privilege.

/CLUSTER

Issues the SET command to each OpenVMS Registry server in the cluster, setting the OpenVMS Registry master server last.

/NODE=(node,...)

Issues the SET command to the OpenVMS Registry servers on the specified nodes, in the order in which you enter the node names. The node names must be in the current cluster.

/START[=(node,...)]

Starts the OpenVMS Registry server on the specified node or nodes in the cluster.

/EXIT[=(node,...)]

Stops the OpenVMS Registry server on the specified node or nodes in the cluster.

/ABORT[=(node,...)]

Aborts the OpenVMS Registry server on the specified node or nodes in the cluster.

/[NO]LOG

Creates a new OpenVMS Registry log file in SYS\$REGISTRY. NOLOG is the default.

D.6 OpenVMS Registry Failover in a Cluster

To increase the availability and reliability of the OpenVMS Registry, you can run multiple OpenVMS Registry servers in a cluster, up to one per node. No matter how many OpenVMS Registry servers you run, you have only one OpenVMS Registry database.

When you run more than one OpenVMS Registry server in a cluster, only one OpenVMS Registry server process is active and writing to the OpenVMS Registry database. The other OpenVMS Registry server processes are standing by.

By default, the first OpenVMS Registry server process that is active in the cluster remains active until either the process no longer exists or the priority among OpenVMS Registry server processes changes.

D.6.1 Changing the Priority of OpenVMS Registry Server Processes

You can change the priority of OpenVMS Registry server processes by creating and modifying the priority value of each node in the cluster that will run the OpenVMS Registry server process: the higher the value, the higher the priority.

Example D-1 shows priority values being assigned so that NODENAME1 will be the active OpenVMS Registry server process in the cluster.

Example D-1 Setting Priority Values

```
$ mcr reg$cp
REG> CREATE VALUE HKEY_LOCAL_MACHINE\SYSTEM\REGISTRY\PRIORITY -
_REG> /NAME=NODENAME1/DATA=15/TYPE=DWORD
_REG> CREATE VALUE HKEY_LOCAL_MACHINE\SYSTEM\REGISTRY\PRIORITY -
_REG> /NAME=NODENAME2/DATA=10/TYPE=DWORD
_REG> CREATE VALUE HKEY_LOCAL_MACHINE\SYSTEM\REGISTRY\PRIORITY -
_REG> /NAME=NODENAME3/DATA=5/TYPE=DWORD
```

In Example D-1, if NODENAME1 shuts down, control of the OpenVMS Registry database passes to the server process on NODENAME2.

Example D-2 shows the system manager increasing the priority value of NODENAME3 to 20.

Example D-2 Changing Priority Values

```
$ mcr reg$cp
REG> MODIFY VALUE HKEY_LOCAL_MACHINE\SYSTEM\REGISTRY\PRIORITY -
_REG> /NAME=NODENAME3/DATA=20/TYPE=DWORD
```

In Example D-2, the OpenVMS Registry server process on NODENAME1 goes into standby mode and the OpenVMS Registry server process on NODENAME3 becomes active.

D.7 Connecting to the OpenVMS Registry from a Windows NT System

To connect to the OpenVMS Registry from a Windows NT system, you must do the following:

- On the OpenVMS system:
 - Install the Advanced Server for OpenVMS.

Managing the OpenVMS Registry

D.7 Connecting to the OpenVMS Registry from a Windows NT System

- Configure the Advanced Server for OpenVMS.
- On the Windows NT system:
 - Install and configure any required hardware.
 - Install and configure the Windows NT Server or Workstation software.

When you access the OpenVMS Registry database from a Windows system, you will have all the privileges granted on your Windows NT system. For example, if you are logged on to the Windows NT system as an Administrator, you will be able to read and write to all keys and values in the OpenVMS Registry. Access to OpenVMS Registry keys is based on your Windows NT user profile (username and Group membership). Connect to the OpenVMS Registry through Advanced Server for OpenVMS; use the Windows Regedt32 application to view and change keys, values, and security settings.

Caution

Be careful when you modify OpenVMS Registry database keys and values. If you damage the OpenVMS Registry database, you can affect all applications and users on the entire OpenVMS system or cluster.

D.8 OpenVMS Registry Quotas

A quota mechanism limits the size of the OpenVMS Registry database. The system assigns a quota to the root key datafile for every OpenVMS Registry file. By default, these root keys are the USERS key (REGISTRY\$USERS.REG) and the LOCAL_MACHINE key (REGISTRY\$LOCAL_MACHINE.REG).

The quota limits the size of the information contained within the file but does not include the size of information stored in other files, even if the files are part of the subtree.

The default quota and file-specific quotas are stored in the OpenVMS Registry under the HKEY_LOCAL_MACHINE\SYSTEM\Registry key.

D.9 OpenVMS Registry Security

A user can access (read and modify) the OpenVMS Registry directly in the following ways:

- From a Windows NT system (through a connection through Advanced Server for OpenVMS)
- Using the OpenVMS Registry system services (\$REGISTRY[W])
- Using the OpenVMS Registry server management utility (REG\$CP)

You can change a key's security attributes only from a Windows NT system—you cannot change a key's security attributes from an OpenVMS system. OpenVMS does not create or manage Windows NT security attributes.

D.10 Backing Up and Restoring the OpenVMS Registry Database

The OpenVMS Registry includes a server management utility that allows you to back up and restore the entire OpenVMS Registry database to or from a file from the OpenVMS DCL prompt as long as you have the required system privileges.

For more information about backing up and restoring the OpenVMS Registry database,

D.11 Using the OpenVMS Registry in an OpenVMS Alpha Mixed-Version Cluster

The OpenVMS Registry Server can run in an OpenVMS Alpha mixed-version cluster. That is, the OpenVMS Registry can run in a cluster that includes OpenVMS versions other than OpenVMS Version 7.2-1; but the OpenVMS Registry server must be running on the node that is running OpenVMS Version 7.2-1.

D.12 Internationalization and Unicode Support

To integrate with Windows NT, the OpenVMS Registry is Unicode compliant. For more information about Unicode, see the *OpenVMS Version 7.2 New Features Manual*.

Glossary

An alphabetical list of terms used in the *OpenVMS System Manager's Manual* and their definitions follow.

access control list (ACL)

A protection mechanism using a more refined level of protection than that available with **UIC-based protection**. ACLs can be used to grant or deny access to individual users or groups of users.

access mode

Any of the four processor access modes in which software executes. Processor access modes prevent system software from inadvertently performing operations that might damage the system. Processor access modes are in order from most to least privileged and protected: kernel, executive, supervisor, and user. When the processor is in any mode other than kernel mode, the processor is inhibited from executing privileged instructions.

account

Each system user has an account. When you log in, you log in under a particular account name and number. This number informs the system where your files are and what kind of access to other files and system facilities you should be given.

accounting files

Files where the system stores information about resource use. Compare with **current accounting file**.

active set

In a multiprocessing system, the subset of processors that have successfully run power-on diagnostics and are actively participating in system operations. Compare with **available set**.

active values

With system parameters, the set of values that is stored in memory and is used by the active system. When the system boots, it reads into memory the **current values** stored in a parameter file on disk.

adjacent node

In a network, a node that is connected to your node by a single physical line.

allocation class

In an OpenVMS Cluster environment, for devices that are dual-ported between two computers, a numeric value used to create a unique, path-independent device name.

answer file

A file in the form SYS\$UPDATE:*product*.ANS. The file is created when you install a product initially, and you specify the Auto-Answer option. The file contains a record of the answers you entered when you ran VMSINSTAL.COM to install that product initially.

application service

A **LAT service** in which LAN users can access only a specific program. Contrast with **general timesharing service**.

area router

In a network, a node that performs routing operations between areas and within its own area. Also called a **level 2 router**. Compare with **level 1 router**.

autoconfiguration

Autoconfiguration is the process of discovering the hardware devices on a system and loading the appropriate device drivers.

autostart feature

A feature that simplifies startup and ensures high availability of execution queues in an OpenVMS Cluster environment. It lets you perform the following actions:

- Start all **autostart queues** on a node with a single command
- Specify a list of nodes (within an OpenVMS Cluster environment) to which a queue can automatically fail over if necessary.

autostart queue

An execution queue that takes advantage of the **autostart feature**. When you create a queue, you can designate it as an autostart queue.

available set

In a multiprocessing system, those processors that have successfully completed the system's power-on hardware diagnostics and may or may not be actively involved in the system. Compare with **active set**.

backlink

In Files-11 disk structure, a pointer to the directory in which a file resides.

banner page

A specially formatted page that prints at the beginning and end of print jobs and files within print jobs. These pages are helpful in identifying and separating output jobs, and the files within those jobs, when they are printed.

base process priority

A base priority value that the system uses to schedule a process. Priorities range from a low of 0 to a high of 31; 0 through 15 are timesharing priorities and 16 through 31 are real-time priorities. Compare with **job scheduling priority**.

batch execution queue

An execution queue that can accept only batch jobs.

batch job

A detached process that sequentially runs one or more command procedures. The user defines the list of command procedures when submitting the job to a batch queue.

batch mode

An execution mode in which you can execute a command procedure by submitting the procedure to a batch queue. When resources are available, the system creates a detached process to execute the commands in the procedure. Usually, processes running in batch mode execute at a lower process priority, to avoid competing with interactive users for system resources.

beginning-of-tape (BOT) marker

A piece of photoreflexive tape that delimits the beginning of the writable area on a tape volume.

binding

On an InfoServer system, a function that creates a **virtual device unit** on a local OpenVMS system.

block

On Files-11 disks, the basic unit by which disk space is allocated (512 8-bit bytes). On magnetic tape, the size of a block is determined by the user.

boot block

Virtual block 1 of the index file. The boot (or bootstrap) block is almost always mapped to the logical block 0 of the volume. The boot block contains the location and size of the **primary bootstrap image**, which is used to boot the system. Certain processors, in order to boot, must read the boot block to obtain the location of the primary bootstrap image.

booting

Also called **bootstrapping**, the process of loading system software from the system disk into processor memory. You must install the operating system before you boot the system for the first time. See also **conversational boot** and **nonstop boot**.

bootstrapping

See **booting**.

bpi

Bits per inch; a measure used for characters of data on tape. Also called **density**.

caching

A performance enhancement in which the system stores information in memory; this includes information about a disk volume's free space, file identifications, quota file entries, and file headers.

capability

On VAX systems, software that makes the services of the vector processor available to system users.

cathedral windows

Segmented windows created from mapping windows; useful for reducing the overhead required to read large files. The Buffered I/O Byte Count Limit (BITLM) limits the number of cathedral windows a user can create.

circuit

In a network, a communications data path that connects **adjacent nodes**. A circuit is not a physical data path but, rather, a logical connection that operates over a physical connection (a line). All input and output (I/O) between nodes takes place over circuits.

cluster

On Files-11 media, a logical grouping of blocks; the basic unit by which disk space is allocated.

See also **OpenVMS Cluster system**.

CMIP

See **Common Management Information Protocol**.

command procedure

A file containing DCL commands and, optionally, data used by those commands. When you execute a command procedure, the system reads the file and executes the commands it contains. This eliminates the need for you to enter each command separately. You can use command procedures to efficiently perform routine tasks. A command procedure can also be executed in **batch mode**.

command string

The complete specification of a command, including the command name, command qualifiers, parameters, and parameter qualifiers. Because a command can be continued on more than one line, the term is used to define the entire command.

Common Management Information Protocol (CMIP)

Draft standard for network management operations on which DECnet-Plus network management protocol is based. CMIP is used for encoding network management operations that can be performed on an entity. CMIP permits the exchange of information between a director and an agent. CMIP supersedes the Phase IV Network Information and Control Exchange (NICE) protocol.

Compact Disc Read-Only Memory (CD-ROM)

Computer discs similar to the CD-ROMs used for audio applications. The major difference is that CD-ROM computer disc players have a digital (rather than an audio) interface.

configuration database

In a network, each node has a configuration database that includes information about the node and other nodes with which it can communicate. The configuration database is made up of a **permanent database** and **volatile database**.

connection manager

In an OpenVMS Cluster environment, the component that dynamically defines the OpenVMS Cluster system and coordinates participation of computers in the cluster.

conversational boot

A booting operation in which you stop to perform special operations—for example, to change system parameter values—before booting. Contrast with **nonstop boot**.

Conversational boot operations are common in programming research and development environments where you must alter operating conditions for experimentation, testing, and debugging.

Coordinated Universal Time (UTC)

International standard for measuring time of day.

crash dump

When the operating system detects an unrecoverable error or an inconsistency within itself that causes the system to fail, it writes the contents of the error log buffers, processor registers, and memory into the **system dump file**.

crash history file

A file storing information about system crashes. Use the Crash Log Utility Extractor (CLUE) to display the contents of the crash history file to understand and resolve the issues responsible for crashes, and to obtain other useful data.

current accounting file

In an OpenVMS Cluster environment, an **accounting file** for a particular node. By default, the current accounting file is SYS\$MANAGER:ACCOUNTNG.DAT.

current values

With system parameters, the set of values that is stored in the default parameter file on disk and are used to boot the system. When the system boots, it reads the current parameter values into memory to create **active values**.

cylinder

On a disk, consists of all **tracks** at the same radius on all recording surfaces of the disk.

data area

One of two divisions of CD-ROM volume space; includes the remaining volume space, beginning with logical sector 16.

data record storage

Within tape files, data records are stored in variable-size data blocks. Each block contains one or more records. RMS provides management of records.

DECevent

On Alpha systems, the event management utility that provides an interface between a system user and the operating system's event log files.

DECnet

Generic name for the software and hardware products that allow various Compaq operating systems to participate in a network. DECnet allows a system to function as a node in a network. See also **DECnet Phase IV**, **DECnet-Plus for OpenVMS**, and **Compaq TCP/IP Services for OpenVMS**.

DECnet Phase IV

Older version of DECnet; available for separate installation as a layered product.

DECnet-Plus for OpenVMS

The version of DECnet that is available with the installation procedure for the OpenVMS operating system. DECnet-Plus is Compaq's implementation of the Digital Network Architecture (DNA) Phase V.

default values

With system parameters, the set of values provided on your distribution kit and stored in the default list. These values allow you to boot any supported configuration.

density

A measurement, in bits per inch, used for characters of data on tape.

device

Hardware that allows access to storage media; also called **drive**.

device control library

A text library that contains user-written modules consisting of text or escape sequences. See also **device control module**.

device control library module

A user-written module in a **device control library**. Device control library modules can be used for the following purposes:

- With programmable printers, to insert device-dependent escape sequences that set up a printer for selected print options such as point size, character set, and bold or italic print.
- With both programmable and nonprogrammable printers, to insert text at specific points in the processing of a print job.

See also **page setup module**, **reset module**, and **setup module**.

device driver

A system component that controls I/O operations for a particular device type. For a device to function on a system, the device must be connected and the device driver must be loaded into memory.

Compaq TCP/IP Services for OpenVMS

Compaq implementation of the TCP/IP protocol suite and Internet services for OpenVMS Alpha and VAX operating systems.)

disk

Physical media on which files reside.

disk quota

A method for maintaining and enforcing limits on the amount of disk space available to users on a public volume. See also **quota file**.

drive

Hardware that allows access to storage media; also called **device**.

end node

In a network, a node that does not perform routing operations.

end-of-tape (EOT) marker

A piece of photorefective tape that delimits the end of the writable area on a tape volume.

ERRFMT process

System process that periodically empties the error log buffers, transforms the descriptions of the errors into standard formats, and stores the formatted information in the **error log file** on the system disk.

error log dump file

On Alpha systems, file to which the contents of error log buffers are written (in addition to the system dump file). The error log dump file is provided so that the system can be updated on reboot to include error log entries that were created but not written at the time of a system crash.

error log file

The operating system automatically records device and CPU error messages in this file. The Error Log utility invokes the **Error Log Report Formatter (ERF)** to selectively report the contents of an error log file.

Error Log Report Formatter (ERF)

A system component invoked by the Error Log utility to selectively report the contents of the **error log file**.

Ethernet

A single shared network channel, with all nodes having equal access to the channel. Ethernet offers local and remote connections as one integral network.

event classes

Categories of security-relevant events. The system always audits several event classes.

executable image

An image that can be run in a process. It is linked with the /EXECUTABLE qualifier (or without the /SHAREABLE qualifier) of the Linker utility.

execution queue

A queue that accepts batch or print jobs for processing. Compare with **generic queue**.

executive

A set of programs in the operating system that controls the running of routines that perform I/O, resource allocation, and program execution. See also **executive routines**.

executive mode

The second most privileged processor **access mode**. OpenVMS Record Management Services (RMS) and many system service procedures execute in executive mode.

executive routines

System routines that detect errors and events and write relevant information into error log buffers in memory. See also **executive**.

expiration date

The Files-11 On-Disk Structure uses the expiration date of a file to track the use of a file. The expiration date aids in the disposal of seldom-used files.

extent

On Files-11 volumes, contiguous blocks allocated to a particular file.

feedback

Information, continuously collected by the **executive**, about the amount of various resources the system uses to process its work load. When run in feedback mode, AUTOGEN analyzes this information and adjusts the values for any related system parameters.

field

In a UAF record, a portion of the record you modify with the Authorize utility. The values you assign to each field perform the following functions:

- Identify the user
- Define the user's work environment
- Control use of system resources

file

On Files-11 media, an array of consecutive virtual blocks, numbered 1 to n , plus a set of attributes with values. A file is either a data file or a directory file. Directories can contain both data files and directory files.

file banner page

A **banner page** that separates files within a job; users can override the file banner page settings you set for a queue.

file header

On a Files-11 volume, describes a portion of a file on the volume. File headers contain information such as the **owner UIC**, **protection code**, creation date and time, and **access control list (ACL)**.

file operation

In the Backup utility, an operation that processes individual files or directories.

Files-11 On-Disk Structure

A logical structure given to information stored on a disk; it is a hierarchical organization of files, their data, and the directories needed to gain access to them.

Files-11 volume

A disk volume that uses Files-11 On-Disk Structure and is mounted on a device.

full backup

See **image backup**.

full names

On VAX systems, hierarchically structured DECnet-Plus node names that can be stored in a DECdns naming service. Full names on VAX systems can be a maximum of 255 bytes long.

gateway

In a TCP/IP network, a computer that is connected to two networks and routes packets from one to the other.

general timesharing service

A LAT service offering processing resources to users in the LAN. Contrast with **application service**.

generic batch queue

A generic queue that can direct jobs only to batch execution queues.

Generic batch queues are typically used in OpenVMS Cluster environments to distribute the batch work load across several nodes.

generic output queue

A generic queue can direct jobs to any output execution queue. Generic output queues are typically used to distribute the output work load among several identical printers.

generic queue

A queue that holds batch or print jobs until they are transferred to an **execution queue** for processing.

A generic queue holds a job until an appropriate execution queue becomes available to initiate the job. The **queue manager** then requeues the job to the available execution queue.

group volume

A volume available to all the users in a group. Compare to **system volume**.

header labels

On magnetic tape, labels containing information such as the file name, creation date, and expiration date. When you create a file on magnetic tape, the magnetic tape file system writes header labels immediately preceding the data block. To access a file on magnetic tape by the file name, the file system searches the tape for the header label set that contains the specified file name.

header resident image

A **known image** for which the header of the image file remains permanently resident in memory, saving one disk I/O operation per file access.

home block

A block in a Files-11 volume that identifies it as a Files-11 volume. Usually, the home block is the next block after the **boot block** (block 0). If for some reason the home block cannot be read (is physically unusable), an alternative block is selected for use as the home block. This block provides specific information about the volume and default values for files on the volume.

host

In a TCP/IP network, each end system connected to a TCP/IP network. Each host has a unique name and address. The local host is the system you are using, and the remote host is the system with which you are communicating.

identification record

A record of a **file header** that contains a summary of disk and volume characteristics.

IDP

See **new domain part**.

image

A collection of procedures and data bound together by the Linker utility to form an executable program. Executable programs can be executed (or run) by a process. Usually, executable programs have the file type .EXE.

image backup

Also called a full backup. A Backup utility operation that saves a copy of all the files on a disk (or volume) to a special file called a **save set**. See also **image operation**.

image compare

A Backup utility operation that compares the contents of entire volumes.

image copy

A Backup utility operation that creates a new Files-11 On-Disk Structure on the output disk and copies an entire volume; the image backup is a logical duplicate of the contents of the disk.

image operation

A Backup utility operation that processes all files on the input disk.

image registry

A file associated with the Image Registry facility. To continue using a compatible application image that depends on a previous operating system version, you can register the image in the Image Registry.

image restore

A Backup utility operation that initializes the output disk and restores an entire volume.

incremental backup

A Backup utility operation that saves only those files that have been created or modified since the most recent backup that was performed using the /RECORD qualifier. (The /RECORD qualifier records the date and time that the files are backed up.)

incremental restore

A Backup utility operation that restores an incremental **save set**.

InfoServer system

An Ethernet-based, high-performance, **virtual device server**. The InfoServer system can serve physical device media and sets of logical disk blocks to client systems in a local area network (LAN). Systems running the appropriate client software can connect to virtual devices served by the InfoServer system and use them as though they are locally attached devices.

initialization file

In certain utilities, a file used each time you invoke the utility. In the initialization file, you can perform tasks such as defining keys and setting up your environment.

installation procedure

The procedure for installing the operating system for the first time. Also, a procedure for installing a layered product.

IRG (interrecord gap)

On magnetic tape, the interval of space between blocks.

job banner pages

banner pages that identify jobs; users cannot override job banner pages that you set for a queue. Compare with **file banner pages**.

job controller

The system process that creates a process to perform the tasks in a batch job.

job scheduling priority

A priority value that the system uses to schedule batch or print jobs in a queue. Job scheduling priorities range from a low of 0 to a high of 255. Compare with **base process priority**.

kernel mode

The most privileged processor **access mode**. The operating system's most privileged services, such as I/O drivers and the pager, run in kernel mode. When in kernel mode, the processor has complete control of, and responsibility for, the system.

key processes

Processes that are dumped immediately following PT, S0/S1, and S2, including transition pages that link back. The system manager can designate additional processes to be treated as key processes. Key processes have priority over other processes in a dump, thus ensuring that the selected processes are successfully written when the dump file is too small to contain all processes.

known file list

An internal data structure on which the system defines **known images**. Each entry in the known file list identifies the file name of the known image and the attributes with which it was installed.

known image

An image installed with the Install utility (INSTALL). When you install an image, the image is assigned attributes and becomes known to the system.

LASTport protocol

A specialized LAN transport protocol, implemented by the InfoServer software, that allows many clients to access InfoServer systems and perform reliable device read and write operations.

The LASTport/DISK protocol and LASTport/TAPE protocol are specialized disk and tape protocols that use the LASTport protocol.

See also **InfoServer system**.

LAT protocol

Protocol, implemented by the LAT software, that allows the operating system to offer resources, or LAT services that **terminal servers** can access.

LAT service announcements

Multicast messages sent by **LAT service nodes** and used to create a database of service nodes available.

LAT service node

A system that supports incoming LAT connections or a system that offers **LAT services**.

LAT services

Computing resources made available to users in the LAN through the LAT software. A LAT service can be a **general timesharing service** or an **application service**.

level 1 router

In a network, a node that performs routing operations within a single area. Compare with **level 2 router**.

level 2 router

In a network, a node that performs routing operations between areas and within its own area. Also called an **area router**. Compare with **level 1 router**.

license

Many software vendors provide software to their customers under an agreement called a license. Although the term *license* can have specific legal connotations, for the purpose of this manual a license refers to the authorization you have to use a product.

The License Management facility (LMF) lets you register, manage, and track software licenses on line. See also **Product Authorization Key (PAK)**.

lines

In a network, physical data paths that connect **adjacent nodes**. Communications lines connect your computer to the DECnet network.

In a TCP/IP network, a line is the physical path over which data can pass from one host to another.

load address

The location in memory (specified in hexadecimal notation) to which the system loads the **bootstrap image**.

Local Area VAXcluster configuration

A VAXcluster configuration in which a single VAX computer serves as the management center of the cluster, plus one or more VAX computers that are connected to this hub.

local cluster

In the System Management utility (SYSMAN), the node from which you are executing SYSMAN.

local node

In a network, the node on which you are working.

In the System Management utility (SYSMAN), the node on which you execute SYSMAN.

Contrast with **remote node**.

logical block

Organizational unit of volume space.

logical block numbering

Begins with the first byte in the volume space and continues in a sequentially ascending order through the remainder of the volume space.

logical link

In a network, connects two processes and carries a stream of two-way communications traffic between the processes over a **circuit**. A single circuit established between two nodes can support many logical links concurrently.

logical name table

A table containing definitions of systemwide logical names that can be used by any process in the system.

logical queue

A special type of generic output queue that transfers print jobs to another output execution queue. You might use this kind of queue to temporarily redirect a queue when the device on which it runs is broken.

logical sector

Organizational unit of a volume; consists of one or more physical sectors. No more than one logical sector can begin in any physical sector.

Logical sectors are numbered in ascending order, with 0 assigned to the logical sector having the lowest physical address containing recorded data. Each logical sector includes a data field made up of 2048 or more bytes (the number of bytes always equals a power of 2).

login command procedure

A command procedure that executes each time a user logs in. Add commands to a login command procedure to execute commands when a user logs in, for example, to set up the user environment.

login (LGI) system parameters

System parameters that control login functions. The names of these system parameters begin with LGI.

loopback tests

In a network, a series of tests to help determine whether the network is operating properly.

lost file

A file that is not linked to a directory. When you delete a directory file (a file with the file type .DIR) without first deleting its subordinate files, the files referred to by that directory become lost files. Lost files are a nonproductive use of disk space and act as debits against a user's **disk quota**.

Magnetic Tape Ancillary Control Process (MTACP)

The internal software process of the operating system that interprets the logical format of standard labeled tape volumes.

maintenance release

A release of the operating system that is applied with an **update procedure**.

managers

Devices on the network through which management is done using the Extensible Simple Network Management Protocol (eSNMP). Managers exchange information with **master agents** or subagents, which are devices such as routers and servers on the network being managed.

mandatory update

A software update that is required immediately after upgrading or installing the operating system.

mass storage control protocol (MSCP) server

In an OpenVMS Cluster environment, the component that implements the MSCP protocol, which is used to communicate with a controller for DSA disks, such as RA-series disks. In conjunction with one or both of the disk class **device drivers** (DUDRIVER, DSDRIVER), the MSCP server implements this protocol on a computer, allowing the computer to function as a storage controller.

master agents

Devices such as routers and servers on the network being managed (using the Extensible Simple Network Management Protocol (eSNMP)). Master agents or subagents exchange information with **managers**, which are the devices on the network through which the management is done.

master file directory (MFD)

The file that contains the name of all user file directories on a disk.

media

The physical substance on which you can store data.

mount verification

A recovery mechanism for disk and tape operations. If a device goes off line or is **write-locked** while mount verification is enabled, you can correct the problem and continue the operation.

multivolume file

A file that is continued on another volume when the data blocks of a file or related files do not physically fit on one volume (a reel of magnetic tape).

network

A means of connecting computers that allows them to share or transfer information or communications. A network includes two or more computers that are connected, and the hardware and software that makes those connections.

network proxy account

A user account that allows users on a **remote node** in a network to access data by way of a local account on your system. Proxy accounts are useful when you want to grant one or more users on a remote node access to specific files but you do not want to give them a private account on your system.

new domain part (IDP)

Unique network identifier that allows users on a DECnet-Plus network to communicate with users on other OSI networks, either through electronic mail, EDI, FTAM, VTP, or other internetwork utilities.

node

In a network, a computer system that is connected to another system in a network—by means of cables, telephone lines, microwave and satellite links, for example.

nonlocal cluster

In the System Management utility (SYSMAN), any cluster other than the one from which you are executing SYSMAN.

nonlocal environment

In the System Management utility (SYSMAN), your environment when you are not working on your local node or within your own cluster.

nonstop boot

The most common booting operation. You perform a nonstop boot if you do not want to stop to perform special operations—for example, to change system parameter values—before booting. Contrast with **conversational boot**.

object

In a network, a process to which a **logical link** connects. Some objects are DECnet programs—for example, the Mail object; other objects are user-written programs.

For two programs to communicate over the network, the source program on the local node establishes a **logical link** with the object on the remote node.

OPCOM messages

Messages broadcast by the Operator Communication Manager (OPCOM). These messages are displayed on **operator terminals** and written to the **operator log file**. The messages might be general messages that you send, user requests, operator replies, or system events.

OPCOM process

The system process that manages Operator Communication Manager (OPCOM) operations.

OpenVMS Cluster system

A loosely coupled configuration of two or more computers and storage subsystems. An OpenVMS Cluster system appears as a single system to the user, even though it shares some or all of the system resources. When a group of computers shares resources in an OpenVMS Cluster environment, the storage and computing resources of all the computers are combined, which can increase the processing power.

operator log file

The Operator Communication Manager (OPCOM) records messages in this file. The file is named SYS\$MANAGER:OPERATOR.LOG.

operator terminals

Terminals designated to display messages broadcast by the Operator Communication Manager (OPCOM). Usually, the console terminal (with the device name OPA0:) is the operator terminal. However, you can designate any user terminal as an operator terminal.

output execution queue

A queue that accepts jobs for processing by a **symbiont**. The **queue manager** sends the symbiont a list of files, which the user defines when submitting the job. An output symbiont transfers data from a disk to an output device. As the symbiont processes each file, it produces output for the device it controls, such as a printer or a terminal.

owner UIC

Used with **UIC-based protection**, usually the UIC of the person who created a file or volume.

page

A unit used for allocating and deallocating memory.

On VAX systems, a page is 512 bytes.

On Alpha systems, a page can be 8 kilobytes (KB) (8192 bytes), 16KB, 32KB, or 64KB. The initial set of Alpha computers uses a page size of 8192 bytes. Compare with **pagelet**.

page file

In a **paging** operation, the file to which the system writes paged portions of memory. Your distribution kit includes a page file named SYS\$SYSTEM:PAGEFILE.SYS. If necessary, SYS\$SYSTEM:PAGEFILE.SYS can be used in place of the system crash dump file.

pagelet

On Alpha systems, a unit of memory in a 512-byte quantity. One Alpha pagelet is the same size as one VAX page. Also, on an Alpha 8KB computer, 16 Alpha pagelets equal 1 Alpha page.

page setup module

A **device control module** inserted at the beginning of each page of a print job.

paging

A memory management operation to efficiently use the physical memory allotted to a process by moving information between physical memory and files stored on disk. In paging, the system moves infrequently used portions of a process workspace out of physical memory to a file. Compare with **swapping**.

PAK

See **Product Authorization Key (PAK)**.

partition

A logical subset of a read/write disk. A single disk can be subdivided into several partitions, each of which can be used independently. The partitions appear to be whole disks.

permanent database

In a network, a permanent copy of the DECnet **configuration database**. When you start the network, the permanent database provides the initial values for the **volatile database**. Changes remain after the network is shut down, but do not affect the current system.

permanently open image

A **known image** where directory information on the image file remains permanently resident in memory, eliminating the usual directory search required to locate a file.

physical dump

A **crash dump** containing the entire contents of physical memory to the **system dump file**. Compare with **selective dump**.

physical operation

In the Backup utility, an operation that copies, saves, restores, or compares an entire volume by logical blocks, ignoring any file structure.

physical sector

Division of a system or data area; smallest addressable unit on an ISO 9660 CD-ROM.

primary bootstrap image

Program that the boot block points to, which allows access to the system disk by finding the **secondary bootstrap image**, SYSBOOT.EXE, and loading it into memory.

On VAX systems, the primary bootstrap image is VMB.EXE.

On Alpha systems, the primary bootstrap image is APB.EXE.

primary page and swap files

The default **page file** and **swap file** provided with your distribution kit. These files are named SYS\$SYSTEM:PAGEFILE.SYS and SYS\$SYSTEM:SWAPFILE.SYS. Contrast with **secondary page and swap files**.

primary processor

In a multiprocessing system, the processor that is either logically or physically attached to the console device and is the target of the console commands that bootstrap the multiprocessing system. The primary processor is responsible for starting other processors in the multiprocessing system. It also serves as the system timekeeper.

print forms

You can use print forms with output queues to determine certain page formatting attributes (such as margins and page length). In addition, the paper stock specified in a form determines whether a job is printed; if the stock of a job's form does not match the stock of the form mounted on the queue, the job is not printed.

Compaq supplies a default print form named DEFAULT. You can create additional forms if users need help formatting output, or if certain print jobs require special paper.

print job

An entry in an output queue that specifies a file or files to be printed on a printer. The user defines the file or files to be printed when submitting the job. When a printer is available, the **queue manager** sends the file to a **symbiont** for formatting and printing.

printer queue

A type of output execution queue that uses a **symbiont** to direct output to a printer. Compare with **server queue** and **terminal queue**.

priority

See **base process priority** or **job scheduling priority**.

private volume

A file-structured disk volume that contains only private files.

privileged image

A **known image** where increased **privileges** are temporarily assigned to any process running the image, permitting the process to exceed its user authorization file (UAF) privilege restrictions during execution of the image. In this way, users with normal privileges can run programs that require higher-than-normal privileges.

privileges

A means of restricting the functions users are authorized to perform on the system. System managers require privileges that are denied to most users.

process limits and quotas

User authorization file (UAF) parameters you can set for a user account to control the usage of system resources by processes in that account. (UAF parameters are different than system parameters.) You set values for process limits and quotas using the Authorize utility.

Product Authorization Key (PAK)

Information, typically on a piece of paper, provided for many Compaq products. The data provided in the PAK allows you to register a software **license** in the license database on a system.

product configuration file (PCF)

Optional POLYCENTER Software Installation utility file that might be supplied by the software manufacturer, or you can create it. A PCF contains responses to some or all of the installation questions for a product. It can provide default or required choices, which might differ from the default choices provided in the PDF.

product database (PDB)

Database that is created automatically by the POLYCENTER Software Installation utility. When products are installed, the files and other objects that make up the product, such as directories and accounts, are recorded in the PDB. The configuration choices made during installation are also recorded.

product description file (PDF)

File provided by the software manufacturer containing all the information the POLYCENTER Software Installation utility needs for installing either a software product or a set of software products.

product text file (PTF)

POLYCENTER Software Installation utility file that is optionally supplied by the software manufacturer. It provides information about the product including product name, producer, configuration choice descriptions, and message text used during product installation.

protected image

A **known image** that is a **shareable image** and contains protected code. Protected code is code that runs in **kernel mode** or **executive mode** but that can be called by a **user mode** image.

protection code

Used with **UIC-based protection**, indicates who is allowed access and for what purposes.

public volume

A Files-11 volume that any user on the system can access and that can contain both private and public files.

queue

Allows users to submit requests for printing or batch processing. The system prints users' print jobs or processes users' batch jobs as resources allow.

queue characteristics

Characteristics you can define and assign to a queue to control the batch or print jobs that execute on the queue.

queue database

A file or files that store information about queues and batch and print jobs.

queue manager

The system component that controls queue activity.

quota file

On Files-11 volumes, the file that records all users who are allowed to use a disk and that shows their current disk usage and their maximum disk allocation. A quota file, QUOTA.SYS, which is stored in directory [000000] with other system files, requires 1 block of disk storage for every 16 entries. See also **disk quotas**.

reconfiguration

After installation of a product, changing the configuration choices made during the installation.

record blocking

On Files-11 volumes, the grouping of individual records into a block, thereby reducing wasted space.

remote node

In a network, a node that is accessible to the node you are working on (the local node) over the network.

In the System Management utility (SYSMAN), any node other than the one on which you are executing SYSMAN.

Contrast with **local node**.

reset module

A **device control module** inserted at the end of each print job. Use reset modules to reset a printer at the end of a job.

resident image

On Alpha systems, a **known image** that improves the performance of a **shareable image**. With a resident image, portions of images that contain code are moved into system space, where they reside on a large single page, thus improving performance.

root volume

The first volume in a **volume set**. Each volume in the volume set is identified by a volume number relative to the root volume, which is always relative to volume 1.

router

In a network, a node that performs routing operations.

routing

In a network of more than two nodes, the process of directing a data message from a source node to a destination node (known as an **end node**). Both routers and end nodes can send messages to and receive messages from other nodes in the network.

In a TCP/IP network, routing is the process of directing a data message from a source host to a destination host. See also **host**.

ruleset

Software routine or function that is analogous to an executable file; used by DECevent.

satellite

In an OpenVMS Cluster system, a computer without a local system disk. A satellite uses disks and tapes that are locally connected to another computer in the cluster. Disks and tapes are served to satellites by means of the mass storage control protocol (MSCP) and the tape mass storage control protocol (TMSCP), respectively.

save set

A special file used by the Backup utility. The Backup utility saves files to a save set and restores files from a save set. Installation and upgrade procedures restore product files from a save set to your system disk.

scalar

A single data item, having one value. Compare with **vector**.

secondary bootstrap image

Image that allows access to the system disk: SYS\$SYSTEM:SYSBOOT.EXE.

secondary page and swap files

Additional **page files** and **swap files** that you might create for performance or disk space reasons. The system uses the space in the secondary files for paging and swapping in addition to the space in the **primary page and swap files**.

secondary processor

In a multiprocessing system, any processor that is not a **primary processor**.

sector

The smallest unit discernible to the Files-11 On-Disk structure. For most Files-11 disks, a sector is equivalent to a block (512 bytes).

On ISO 9660 volumes, a uniquely addressable unit; each sector on a CD-ROM comprises a sequence of 2048 8-bit bytes.

security audit log file

A clusterwide file that contains a record of security events on the system. Using the ANALYZE/AUDIT command, you can produce reports and summaries of security events from the security audit log file.

selective dump

A **crash dump** containing only those portions of memory most likely to be useful in a crash dump analysis. A selective dump is useful when sufficient disk space is not available to hold all physical memory. Compare with **physical dump**.

selective operation

A Backup utility operation that processes files or volumes selectively, according to criteria such as version number, file type, UIC, date and time of creation, expiration date, or modification date.

sequential organization

On magnetic tape media, the organization of data; that is, data is organized in the order in which it is written to the tape.

server queue

A type of output execution queue that uses a user-modified or user-written **symbiont** to process the files that belong to print jobs in the queue. Compare with **printer queue** and **terminal queue**.

setup module

A **device control module** inserted at the beginning of a file in a print job.

shareable image

An image linked with the /SHAREABLE qualifier of the Linker utility; it must subsequently be linked into an executable image to be used. Shareable images are sometimes referred to as **linkable images**.

shared image

A **known image** for which more than one user can access the read-only and non-copy-on-reference read/write sections of the image concurrently, so that only one copy of those sections ever needs to be in physical memory.

shared resource

In an OpenVMS Cluster environment, a resource (such as a disk or a queue) that any node in the cluster can access. Data files, application programs, and printers are some items that can be accessed by users on a cluster with shared resources, without regard to the particular node on which the files or program or printer might physically reside.

site-independent startup command procedure

A command procedure that executes each time a system boots, and manages startup of a system. This file, named SYS\$STARTUP:STARTUP.COM, is required on all systems, regardless of site-specific requirements. *Do not modify this file.* Compare with **site-specific startup command procedure**.

site-specific startup command procedure

A command procedure that executes each time a system boots. Unlike the **site-independent startup command procedure**, you can add commands to site-specific procedures to perform operations that vary from site to site.

sizing

The process of matching the allocation of system resources (memory and disk space) with the workload requirements of your site. Use the AUTOGEN command procedure to automatically size your system.

slicing

On Alpha systems, a feature that lets the operating system split the contents of images and sort the pieces so that they can be placed with other pieces that have the same page protection in the same area of memory. Consequently, translation buffers on Alpha systems are used more efficiently than if the loadable executive images or the shareable images were loaded in the traditional manner.

source disk

In the command procedures VMSINSTAL.COM or VMSKITBLD.COM, the disk from which you copy files. Compare with **target disk**.

spooled printer

A printer set up to write output to an intermediate storage device (such as a disk). Spool printers if your system runs applications that write or copy data directly to printers rather than submitting print jobs to a queue. In this way, printers remain available to other system users while the program is running.

startup database

A file that contains information used to start up system software. For example, the **site-independent startup command procedure** uses information in a startup database named STARTUP\$STARTUP_VMS to start the operating system. It uses information in a startup database named STARTUP\$STARTUP_LAYERED to start layered products.

swap file

In a **swapping** operation, the file to which the system writes swapped portions of memory. Your distribution kit includes a swap file named SYS\$SYSTEM:SWAPFILE.SYS.

swapping

A memory management operation to efficiently use the physical memory allotted to an entire system by moving information between physical memory and files stored on disk. In swapping, the system moves the entire workspace of a less active process out of physical memory to a file. Compare with **paging**.

symbiont

Used with an output queue, a process for formatting of print jobs and sending them to a printer.

The standard print symbiont provided by the operating system is named PRTSMB and is designed to print files on basic output devices. The LAT print symbiont LATSVM is used to print files on output devices attached to a **terminal server**.

SYSGEN parameters

See **system parameters**.

system area

One of two divisions of CD-ROM volume space; includes logical sectors 0 through 15. Reserved for system use.

System Communications Services (SCS)

In an OpenVMS Cluster environment, software that implements intercomputer communication, according to the Compaq Systems Communications Architecture (SCA).

system disk

Disk on which operating system files are stored.

system dump file

The file into which the operating system writes the contents of the error log buffers, processor registers, and memory when it detects an unrecoverable error or an inconsistency within itself that causes the system to fail. See also **crash dump**.

system image

An image that does not run under the control of the operating system. It is intended for standalone operation only. The content and format of a system image differs from that of a **shareable image** and an **executable image**.

system image snapshot

A record of the system setup used with the Snapshot facility.

system messages

Messages returned by the system when you enter commands in DCL or in utilities. These messages help you understand the result of each command.

system parameters

Parameters for which you can set values to control how the system functions. Values of system parameters control a wide range of system functions including but not limited to memory management, process scheduling, and system security.

system startup

Series of tasks that an operating system does to start up the system.

system volume

A volume available to all the users on a system. Compare to **group volume**.

systemwide logical name

A logical name that applies to the entire system. It is defined in the system logical name table and can be used by any process in a system.

tape mass storage control protocol (TMSCP) server

In an OpenVMS Cluster environment, the component that implements the TMSCP protocol, which is used to communicate with a controller for local MSCP tapes, such as TU-series tapes. In conjunction with the tape class **device driver** (TUDRIVER), the TMSCP server implements this protocol on a processor, allowing the processor to function as a storage controller.

target disk

In VMSINSTAL.COM or VMSKITBLD.COM, the disk to which you move the system files. Compare with **source disk**.

terminal queue

A type of output execution queue that uses a **symbiont** to direct output to a terminal printer. Compare with **printer queue** and **server queue**.

terminal servers

Communication devices dedicated for connecting terminals, modems, or printers to a local area network (LAN) and to other systems within a LAN. See also **LAT protocol**.

time differential factor (TDF)

Difference between local system time and Coordinated Universal Time (UTC). See also **Coordinated Universal Time**.

track

On a disk, the collection of **sectors** (or blocks, on Files-11 volumes) at a single radius on one recording surface of the disk. It is accessible to a given read/write head position on the disk device.

trailer labels

On magnetic tape, labels similar to **header labels**, but which are written following the file.

transaction group

A group of computers involved in DECdtm transactions whose SCSNODE names must obey specific rules.

transaction log

A file that stores information about DECdtm transactions performed on a node. It is of file type .LM\$JOURNAL.

trusted logical names

Logical names associated with **executive mode** or **kernel mode**.

tuning

The process of altering various system values to obtain the optimum *overall* performance possible from any given configuration and work load.

UAF

See **user authorization file (UAF)**.

UETP (User Environment Test Package)

A software package designed to test whether the OpenVMS operating system is installed correctly.

UIC

See **user identification code (UIC)**.

UIC-based protection

A protection mechanism based on the **user identification code (UIC)** and applied to all protected objects. Compare with **access control list (ACL)**.

update procedure

Procedure used if you have a previous version of the operating system and you want to make minor fixes to it. When you update the operating system, the update procedure replaces some system files.

upgrade procedure

If you are already running a standard version of the operating system, you can use the upgrade procedure to obtain a higher version.

user authorization file (UAF)

A file containing an entry for every user that you authorize to gain access to the system. Each entry identifies the user name, password, default account, UIC (user identification code), quotas, limits, and privileges assigned to individuals who use the system.

User Environment Test Package (UETP)

See **UETP**.

user identification code (UIC)

The pair of numbers assigned to users, files, and other system objects, that specify the type of access available to the owner, group, world, and system. The UIC consists of a group number and a member number separated by a comma and enclosed within square brackets. Same as UIC. See also **account** and **UIC-based protection**.

user mode

The least privileged processor **access mode**. User processes and run-time library routines run in user mode.

UTC

See **Coordinated Universal Time**.

utility program

A program supplied by Compaq that performs a set of related operations. For example, the Backup utility (BACKUP) allows you to save and restore files.

VAXport drivers

In a VAXcluster environment, **device drivers** that control the communication paths between local and remote ports. (Examples are PADRIVER for the CI, PEDRIVER for the LAN, and PIDRIVER for the DSSI.)

vector

On VAX systems, a group of related **scalar** values, or elements, all of the same data type.

vector-capable systems

On VAX systems, those systems that comply with the VAX vector architecture.

vector consumer

On VAX systems, a process requiring the vector capability and having a vector context.

vector-present processor

On VAX systems, an integrated scalar-vector processor pair, included in a VAX vector processing system configuration.

virtual device server

Serves physical device media and sets of logical disk blocks to client systems in a local area network (LAN). Systems running the appropriate client software can connect to virtual devices as though they are locally attached devices. A virtual device server does not impose a file system on the virtual devices that it serves. See also **InfoServer system**.

virtual device unit

With an InfoServer system, a virtual device that represents the local OpenVMS context for a volume that resides on a remote server.

Virtual disk units have a device name in the DAD*n*: format. Virtual tape units have a device name in the MAD*n*: format.

See also **binding**, **InfoServer system**, and **virtual device server**.

volatile database

On a node in a network, a working copy of the DECnet **configuration database** that reflects current network conditions. Contrast with **permanent database**.

volume

Disk or tape media that has been prepared for use by creating a new file structure on it and mounting it on a device.

volume set

A collection of disk volumes bound into a single entity by the DCL command MOUNT/BIND. To users, a volume set looks like a single, large volume.

Also, the volumes on which a set of multivolume files is recorded.

volume space

Set of all logical sectors on a volume containing information about the volume.

writable image

A **known image** for which a shared non-copy-on-reference writable section is removed from physical memory (for paging reasons or because no processes are referencing it), and it is written back to the image file.

write lock

A device becomes write-locked when a hardware or user error occurs while a disk or magnetic tape volume is mounted for a write operation. For example, if a disk is write-locked or a tape is missing a write ring, the hardware generates an error.

A

- Aborting job status, *Vol 1*, 14–66
- ACCDATE function, *Vol 1*, 10–18
- Access control entries
 - See ACEs
- Access Control Entry (ACE), *Vol 1*, 10–14
- Access control lists
 - See ACLs
- Access dates, *Vol 1*, 10–18
 - DCL, *Vol 1*, 10–18
- Accessibility field
 - tape file system checks, *Vol 1*, 10–33
- Access modes
 - recommended for logical names of system components, *Vol 1*, 5–9
- Access types
 - abbreviations of, *Vol 1*, 12–8
 - checking when writing files to tape volumes, *Vol 1*, 10–33
 - protecting
 - directory files, *Vol 1*, 10–26
 - disk files, *Vol 1*, 10–21
 - protection codes and, *Vol 1*, 12–8
 - read, *Vol 1*, 10–28
 - continuation volumes, *Vol 1*, 9–56
 - write
 - continuation volumes, *Vol 1*, 9–55
- Account expiration, *Vol 1*, 7–45
- ACCOUNTING command, *Vol 2*, 21–4
- Accounting groups
 - setting up, *Vol 2*, 21–5
- Accounting utility (ACCOUNTING), *Vol 2*, 21–4
- ACCOUNTNG.DAT file, *Vol 2*, 21–3
- ACCOUNTNG logical name, *Vol 2*, 21–4
- Accounts
 - access, *Vol 1*, 7–16
 - adding, *Vol 1*, 7–17, 7–18
 - with ADDUSER.COM, *Vol 1*, 7–18
 - adding proxy logins, *Vol 1*, 7–38
 - automatic login, *Vol 1*, 7–34
 - captive, *Vol 1*, 7–13
 - deleting, *Vol 1*, 7–25
 - directory, *Vol 1*, 7–15
 - disabling, *Vol 1*, 7–27
 - Mail, *Vol 1*, 7–41
 - maintaining, *Vol 1*, 7–24
 - Accounts (cont'd)
 - network proxy, *Vol 1*, 7–37
 - project, *Vol 1*, 7–35
 - restricted, *Vol 1*, 7–13
 - restricting use, *Vol 1*, 7–27
 - security, *Vol 1*, 7–16
 - using ADDUSER.COM, *Vol 1*, 7–18
- ACEs (access control entries)
 - adding to ACL after file is created, *Vol 1*, 10–22
 - Creator ACEs, *Vol 1*, 12–10
 - Identifier ACEs, *Vol 1*, 12–9
 - none for subdirectories, *Vol 1*, 10–27
 - placing on ODS-5 volume to prevent access, *Vol 1*, 10–15
 - security alarm, *Vol 1*, 12–10
 - security audit, *Vol 1*, 12–10
 - subsystem ACEs, *Vol 1*, 12–10
 - using to deny access to ODS-5 volumes, *Vol 1*, 10–14
 - using to override default UIC protection, *Vol 1*, 10–27
- ACL editor
 - invoking, *Vol 1*, 12–12
- ACLs (access control lists), *Vol 1*, 7–16, 7–35
 - default protection, *Vol 1*, 10–27
 - on public volumes, *Vol 1*, 9–21
 - on queues, *Vol 1*, 14–25
 - on vector capability object, *Vol 2*, 28–8
 - SHOW ACL command, *Vol 1*, 10–20
 - using to prevent access to ODS-5 volumes, *Vol 1*, 10–14
- Activating an autostart queue, *Vol 1*, 14–16, 14–17
 - on LAT queues, *Vol 1*, 14–4
 - relationship to starting an autostart queue, *Vol 1*, 14–4
- Active disks
 - backing up, *Vol 1*, 11–66
- Active sets, *Vol 2*, 28–2
 - displaying, *Vol 2*, 28–3
- Active system parameters, *Vol 2*, 15–3, 15–26
- Adding comments to Compaq messages in the Help Message database, *Vol 1*, 5–30
- Adding files to the system disk, *Vol 1*, 5–1
- Adding messages to the Help Message database, *Vol 1*, 5–32

ADDUSER.COM command procedure, *Vol 1*, 7–18
 ADD_DUMPFILE symbol, *Vol 2*, 16–31
 ADD_ERRLOGDUMP symbol, *Vol 2*, 16–31
 ADD_PAGEFILE_n_SIZE symbol, *Vol 2*, 16–31
 ADD_PAGEFILE symbol, *Vol 2*, 16–31
 ADD_ prefix for AUTOGEN, *Vol 2*, 15–19
 ADD_SWAPFILE_n_SIZE symbol, *Vol 2*, 16–31
 ADD_SWAPFILE symbol, *Vol 2*, 16–31
 Advanced Server for OpenVMS, *Vol 2*, D–1, D–10
 environment, *Vol 1*, 10–1
 AGEN\$FEEDBACK.DAT file
 description, *Vol 2*, 15–11
 AGEN\$FEEDBACK_REQ_TIME logical name,
 Vol 2, 15–22
 AGEN\$PARAMS.REPORT
 logs sizes of page, swap, and dump files, *Vol 2*,
 15–18
 AGEN\$PARAMS.REPORT file, *Vol 2*, 15–11
 managing, *Vol 2*, 15–12
 sample, *Vol 2*, 15–11
 Alarms
 security
 enabling, *Vol 2*, 20–30
 security applications, *Vol 1*, 12–13
 ALF (automatic login facility)
 See Automatic login facility
 setting up an automatic login account, *Vol 1*,
 7–34
 Aliases, *Vol 1*, 10–47
 Alias file name
 assigning, *Vol 1*, 10–25
 Alias OpenVMS Cluster node name, *Vol 1*, 2–16
 Aligning preprinted forms, *Vol 1*, 14–70, 14–71
 Aligning queue status, *Vol 1*, 14–59
 Alignment data, *Vol 1*, 14–71
 ALLOCATE command
 for a particular type of device, *Vol 1*, 9–15
 tape drive, *Vol 1*, 10–38
 to allocate a device, *Vol 1*, 9–15
 Allocating
 disk drives, *Vol 1*, 9–15
 allocating a particular type of device, *Vol*
 1, 9–15
 space on disk volume, *Vol 1*, 9–66
 tape drives, *Vol 1*, 9–15
 ALPHAVMSSYS.PAR file, *Vol 1*, 4–2; *Vol 2*, 15–3
 initializing parameters at boot time, *Vol 2*,
 15–36
 Alternate Root
 VMSINSTAL.COM option, *Vol 1*, 3–14
 restriction, *Vol 1*, 3–14
 specifying for software installations, *Vol 1*,
 3–9
 Alternate root directory
 adding to an existing system disk, *Vol 1*, 2–32
 Alternate startup command procedure
 specifying, *Vol 1*, 4–14
 as the default, *Vol 1*, 4–15
 Alternate Working Device
 VMSINSTAL.COM option, *Vol 1*, 3–10
 ANALYZE/AUDIT command
 See also Audit Analysis utility
 generating security reports, *Vol 2*, 20–31
 Analyze/Disk_Structure utility (ANALYZE/DISK_
 STRUCTURE)
 builds BITMAP.SYS file, *Vol 2*, A–8
 checks validity of files, *Vol 2*, A–8
 commands, *Vol 1*, 9–71
 creates version of quota file, *Vol 2*, A–9
 creating disk usage file, *Vol 1*, 9–70
 directing output, *Vol 1*, 9–70
 files used by, *Vol 2*, A–4
 identification record, *Vol 1*, 9–70
 listing file information, *Vol 1*, 9–70
 recovering lost files, *Vol 1*, 9–72
 repairing disk errors, *Vol 1*, 9–71
 reporting disk errors, *Vol 1*, 9–71
 uses of, *Vol 1*, 9–70
 uses VOLSET.SYS to locate volumes in set, *Vol*
 2, A–9
 using after BACKUP/IMAGE operation, *Vol 1*,
 11–4
 verifies files in directory structure, *Vol 2*, A–9
 ANALYZE/ERROR_LOG command, *Vol 2*, 20–5
 Error Log utility, *Vol 2*, 20–7
 excluding unknown entries, *Vol 2*, 20–8
 specifying output, *Vol 2*, 20–7
 ANALYZE/MEDIA command
 to invoke Bad Block Locator utility, *Vol 1*, 9–80
 Announcements
 creating systemwide, *Vol 1*, 5–14
 displaying system, *Vol 1*, 5–15
 Answer file (for software installation), *Vol 1*, 3–10
 APB.EXE file, *Vol 1*, 4–21
 role in boot process, *Vol 1*, 4–2
 Application images
 registering with the Image Registry facility,
 Vol 1, 5–23
 Applications
 untested
 prevented from accessing ODS-5 volumes,
 Vol 1, 10–14
 Archive files
 creating for each node in a cluster, *Vol 2*, 20–28
 Arrow keys
 functions of, *Vol 2*, 22–7
 ASCII character set
 percent sign, *Vol 1*, 10–32
 ASSIGN/MERGE command, *Vol 1*, 14–64
 Assigning
 a default form to a queue, *Vol 1*, 14–50
 a logical queue, *Vol 1*, 14–64
 characteristics to a queue, *Vol 1*, 14–31

Assigning (cont'd)

- libraries to queues, *Vol 1*, 14–54
- reset modules to a queues, *Vol 1*, 14–55
- Assisted merge and volume shadowing, *Vol 1*, 11–45
- Asterisk (*)
 - as wildcard character, *Vol 1*, 10–31
- ASTLM process limit, *Vol 1*, 7–43
 - value for efficient backups, *Vol 1*, 11–11
- Asymmetric vector processing configuration, *Vol 2*, 28–4
- Asynchronous DECnet
 - using virtual terminals, *Vol 1*, 8–22, 8–23
- Attached processors, *Vol 2*, 28–2
- ATTDATE function, *Vol 1*, 10–18
- Audit Analysis utility (ANALYZE/AUDIT), *Vol 1*, 12–14
 - See also ANALYZE/AUDIT command
 - generating security reports, *Vol 2*, 20–31
- Auditing
 - security, *Vol 1*, 12–13
 - See also Security audit log files; Security auditing
 - displaying using SHOW AUDIT command, *Vol 2*, 20–28
- Audit log files, *Vol 2*, 20–2
 - See also Security audit log files
- Audit server processes
 - creation during system startup, *Vol 1*, 5–5
- Authorization files, *Vol 1*, 7–5
- Authorize utility (AUTHORIZE)
 - ADD/IDENTIFIER command, *Vol 1*, 7–35
 - ADD command, *Vol 1*, 7–17
 - adding a user account, *Vol 1*, 7–17
 - checking UAF quotas for software installation, *Vol 1*, 3–3
 - GRANT/IDENTIFIER command, *Vol 1*, 7–35
 - limiting page file usage, *Vol 2*, 16–11
 - listing user records, *Vol 1*, 7–24
 - modifying a user account, *Vol 1*, 7–23
 - modifying process limits for SYSTEM account, *Vol 1*, 3–4
 - reducing process paging, *Vol 2*, 28–7
 - restricting login hours with, *Vol 2*, 17–4
 - restricting system use with, *Vol 2*, 17–4
 - setting process quotas for efficient backups, *Vol 1*, 11–10
- Autoconfiguration
 - See also AUTOCONFIGURE command
 - benefits, *Vol 1*, 8–7
 - DECnet-Plus addresses, *Vol 2*, 23–9
 - definition, *Vol 1*, 5–8, 8–7
 - in system startup, *Vol 1*, 5–4
 - suppressing, *Vol 1*, 5–8, 8–10
 - of devices, *Vol 1*, 8–11

AUTOCONFIGURE command

- See also Autoconfiguration; IO
- AUTOCONFIGURE command
 - in SYSGEN (VAX), *Vol 1*, 8–7
 - in system startup, *Vol 1*, 5–4, 5–7
 - suppressing, *Vol 1*, 5–8, 8–10
- AUTOGEN.COM command procedure, *Vol 2*, 15–8
 - See also AUTOGEN feedback;
 - MODPARAMS.DAT file
 - ADD_ prefix, *Vol 2*, 15–19
 - AGEN\$PARAMS.REPORT file, *Vol 2*, 15–11
 - as recommended method for changing system parameters, *Vol 2*, 15–5
 - automatic management of page, swap, and dump files, *Vol 2*, 16–28
 - calculation of page file size, *Vol 2*, 16–4, 16–28, 16–30, 16–33
 - calculation of swap file size, *Vol 2*, 16–6, 16–28, 16–30, 16–33
 - calculation of system dump file size, *Vol 2*, 16–2, 16–28, 16–30, 16–33
 - changing page, swap, and dump file sizes, *Vol 2*, 16–28
 - changing system parameters, *Vol 2*, 15–3, 15–4
 - controlling operations performed by, *Vol 2*, 15–10
 - controlling system parameter values set by, *Vol 2*, 15–18
 - converting system parameter values for use with, *Vol 2*, 15–5
 - creating page, swap, and dump files, *Vol 2*, 16–31
 - defining the number of VAXcluster nodes for, *Vol 2*, 15–21
 - displaying page, swap, and dump file size calculations, *Vol 2*, 16–29
 - end phase
 - specifying when invoking, *Vol 2*, 15–9
 - executing
 - in batch, *Vol 2*, 15–23
 - interactively, *Vol 2*, 15–18
 - execution mode
 - specifying when invoking, *Vol 2*, 15–9
 - failure executing SYCONFIG.COM, *Vol 1*, 8–11
 - feedback
 - See AUTOGEN feedback
 - functions of, *Vol 2*, 15–9
 - installing page, swap, and dump files, *Vol 2*, 16–29
 - instructions for using, *Vol 2*, 15–18
 - invoking, *Vol 2*, 15–9
 - MODPARAMS.DAT file
 - See MODPARAMS.DAT file
 - parameters to, *Vol 2*, 15–9
 - performance tuning, *Vol 2*, 17–5
 - phases
 - order of, *Vol 2*, 15–16

AUTOGEN.COM command procedure
 phases (cont'd)
 specifying when invoking, *Vol 2*, 15–9
 recommendation
 for using to size page, swap, and dump files, *Vol 2*, 16–28
 restrictions
 for changing file sizes, *Vol 2*, 16–30
 for specifying system page, swap, and dump file sizes, *Vol 2*, 16–30
 reviewing calculations of, *Vol 2*, 15–10, 15–18
 SETPARAMS.DAT file, *Vol 2*, 15–18
 specifying an alternate default startup command procedure, *Vol 1*, 4–15
 specifying location of page, swap, and dump files, *Vol 2*, 16–31
 specifying number of Ethernet adapters in, *Vol 2*, 15–21
 specifying size of page, swap, and dump files of individual files, *Vol 2*, 16–29, 16–31
 of total file space, *Vol 2*, 16–29, 16–31
 specifying values for parameters not calculated by, *Vol 2*, 15–20
 start phase
 specifying when invoking, *Vol 2*, 15–9
 system parameters affected by, *Vol 2*, 15–10
 timing of file size calculations, *Vol 2*, 16–30
 types of data collected by, *Vol 2*, 15–9
 when to run, *Vol 2*, 15–8
 AUTOGEN feedback, *Vol 2*, 15–8 to 15–23
 checks performed on, *Vol 2*, 15–11
 collection of, *Vol 2*, 15–11
 examining effect on parameters, *Vol 2*, 15–11
 file stored in, *Vol 2*, 15–11
 importance of system work load, *Vol 2*, 15–11
 improving system performance, *Vol 2*, 15–10
 maximum age, *Vol 2*, 15–11
 minimum age, *Vol 2*, 15–11, 15–22
 report file, *Vol 2*, 15–11
 sample, *Vol 2*, 15–11
 sending automatically, *Vol 2*, 15–8, 15–23
 resources affected by, *Vol 2*, 15–10
 saving during system shutdown, *Vol 1*, 4–24
 AUTOGEN utility
 Reserved Memory Registry, *Vol 2*, 17–26
 Automatic configurations
 of devices, *Vol 1*, 8–7
 Automatic start
 See also Autostart feature
 of queue manager, *Vol 1*, 13–3, 13–8, 13–10
 Automatic volume switching, *Vol 1*, 9–53
 Autostart feature
 See also Autostart queues
 description, *Vol 1*, 14–4
 disabling, *Vol 1*, 13–9, 14–62
 before shutting down a node, *Vol 1*, 14–63
 enabling, *Vol 1*, 14–4, 14–17
 on LAT queues, *Vol 1*, 14–4, 14–9

Autostart feature (cont'd)
 recommended use, *Vol 1*, 14–9
 Autostart queues
 See also Autostart feature
 activating, *Vol 1*, 14–4, 14–16, 14–17
 activating inactive queue status, *Vol 1*, 14–59
 creating, *Vol 1*, 14–15, 14–17
 preventing from starting, *Vol 1*, 14–62
 recommended use, *Vol 1*, 14–9
 relationship between activating and starting, *Vol 1*, 14–4
 starting, *Vol 1*, 14–17
 in startup command procedure, *Vol 1*, 14–17
 troubleshooting, *Vol 1*, 14–76
 with LAT printers, *Vol 1*, 14–4, 14–9
 AUTO_POSITIONING command
 SHOW CLUSTER, *Vol 2*, 22–12
 Availability
 of devices
 OPCOM message, *Vol 1*, 9–73
 of queue manager, *Vol 1*, 13–9, 13–19
 of queues, *Vol 1*, 14–4, 14–16
 Available queue status, *Vol 1*, 14–59
 Available sets, *Vol 2*, 28–2

B

Backing up
 converting from ODS-2 to ODS-5 volumes, *Vol 1*, 9–37
 performing in command procedures, *Vol 1*, 9–60
 system disk, *Vol 1*, 5–35
 after installation, *Vol 1*, 3–9
 Backing up the OpenVMS Registry, *Vol 2*, D–11
 Backlinks
 definition, *Vol 1*, 9–71
 BACKUP.SYS file
 See Backup log file
 BACKUP /PHYSICAL command
 for backup of ODS-5 volumes on VAX systems, *Vol 1*, 10–9
 BACKUP command
 and save sets, *Vol 1*, 11–25
 /EXACT_ORDER qualifier, *Vol 1*, 11–23
 for backing up directories, *Vol 1*, 11–26
 for copying directories, *Vol 1*, 11–24
 for copying files, *Vol 1*, 11–24
 for image backups, *Vol 1*, 11–33, 11–34
 for incremental backups, *Vol 1*, 11–36, 11–37
 format, *Vol 1*, 11–5
 /GROUP_SIZE qualifier, *Vol 1*, 11–66
 /IGNORE=LABEL_PROCESSING qualifier, *Vol 1*, 11–24, 11–69
 /IGNORE qualifier, *Vol 1*, 11–32, 11–66
 /IMAGE qualifier, *Vol 1*, 11–33, 11–34, 11–59
 /INITIALIZE qualifier, *Vol 1*, 11–15

BACKUP command (cont'd)

- input specifier, *Vol 1*, 11-5
- in VMSINSTAL.COM command procedure, *Vol 1*, 3-9
- /JOURNAL qualifier, *Vol 1*, 11-28
- /LABEL qualifier, *Vol 1*, 11-14, 11-25, 11-36
- /LOG qualifier, *Vol 1*, 11-67
- output specifier, *Vol 1*, 11-5
- /PHYSICAL qualifier, *Vol 1*, 11-59
- qualifiers, *Vol 1*, 11-5
- /RECORD qualifier, *Vol 1*, 11-33, 11-34, 11-36, 11-37
- /REWIND qualifier, *Vol 1*, 11-14, 11-36
- /SAVE_SET qualifier, *Vol 1*, 11-34, 11-37
- /SINCE qualifier, *Vol 1*, 11-36, 11-37
- /VERIFY qualifier, *Vol 1*, 11-67
 - backing up a save set twice, *Vol 1*, 11-67
 - with multiple output devices, *Vol 1*, 11-26, 11-34, 11-35

Backup log file

- BACKUP.SYS, *Vol 2*, A-9
- reserved file, *Vol 2*, A-9

Backup Manager, *Vol 1*, 11-6

- features, *Vol 1*, 11-6
- getting started, *Vol 1*, 11-7
- types of help available, *Vol 1*, 11-6

BACKUP media

- Files-11 disk save set, *Vol 1*, 11-8
- magnetic tape save set, *Vol 1*, 11-7
- network save set, *Vol 1*, 11-8
- sequential-disk save set, *Vol 1*, 11-9

Backups

- image
 - See Image backup
- incremental
 - See Incremental backup
- standalone
 - See Standalone BACKUP

Backup utility (BACKUP)

- backing up active disks, *Vol 1*, 11-66
- backing up a save set twice, *Vol 1*, 11-67
- backing up shadow sets, *Vol 1*, 11-43
- backing up system disks, *Vol 1*, 11-32
- backing up tapes, *Vol 1*, 11-22
- backing up user disks, *Vol 1*, 11-32
- Backup Manager, *Vol 1*, 11-5
- command format, *Vol 1*, 11-5
- command procedures, *Vol 1*, 11-39
- command qualifiers, *Vol 1*, 11-5
- converting from ODS-5 to ODS-2 file names, *Vol 1*, 9-39
- copying dump file, *Vol 2*, 16-5
- copying system dump file, *Vol 2*, 16-23
- copying the queue database, *Vol 1*, 13-13
- data integrity mechanisms, *Vol 1*, 11-65
- file versions, *Vol 1*, 11-24

Backup utility (BACKUP) (cont'd)

- incremental backups using PATHWORKS, *Vol 1*, 11-38
 - InfoServer tapes, *Vol 1*, 11-64
 - limitations with Extended File Specifications, *Vol 1*, 9-10
 - MOUNT messages when backing up tapes, *Vol 1*, 11-22
 - open files during a backup, *Vol 1*, 11-32, 11-66
 - restoring shadow sets, *Vol 1*, 11-50
 - save set, *Vol 1*, 11-25
 - stores save-set file in MFD, *Vol 2*, A-9
 - transferring information, *Vol 1*, 10-34
 - use by VMSINSTAL.COM command procedure, *Vol 1*, 3-9
 - using /PHYSICAL qualifier on VAX systems, *Vol 1*, 10-9
 - using on workstations, *Vol 1*, 11-39
 - using the ANALYZE/DISK_STRUCTURE command, *Vol 1*, 11-4
 - using the menu system, *Vol 1*, 11-32
 - using to back up directories, *Vol 1*, 11-26
 - using to copy directories, *Vol 1*, 11-24
 - using to copy files, *Vol 1*, 11-24
- ## BACKUSER.COM command procedure, *Vol 1*, 11-39
- ## BADBLK.SYS file
- See Bad block file
- ## Bad block file
- BADBLK.SYS, *Vol 2*, A-8
 - definition, *Vol 2*, A-8
 - reserved file, *Vol 2*, A-8
- ## Bad Block Locator utility (BAD)
- analyzing media for bad blocks, *Vol 1*, 9-81
 - detecting media errors, *Vol 1*, 9-80
 - invoking with ANALYZE/MEDIA, *Vol 1*, 9-81
- ## BADLOG.SYS file
- See Pending bad block log file
- ## Banner pages
- commands used with, *Vol 1*, 14-46
 - definition, *Vol 1*, 14-38
 - file, *Vol 1*, 14-46
 - job, *Vol 1*, 14-46
- ## BAP system parameters, *Vol 2*, 15-37
- ## BASEENVIRON phase of system startup, *Vol 1*, 5-4, 5-18
- ## Base priority
- See Priority, base
- ## Batch and print queuing system, *Vol 1*, 13-2
- See also Queue configurations
 - components, *Vol 1*, 13-1
 - in OpenVMS Cluster environments
 - with multiple system disk, *Vol 1*, 13-6
 - PRINT/DELETE command, *Vol 1*, 14-78
 - queue database
 - location of files, *Vol 1*, 13-5
 - mounting disk for, *Vol 1*, 13-6

Batch and print queuing system (cont'd)

- queuing process, *Vol 1, 14-2*
- sample configurations, *Vol 1, 14-5 to 14-13*
- starting in system startup, *Vol 1, 5-12*
- steps for setting up, *Vol 1, 14-13*
- SUBMIT/DELETE command, *Vol 1, 14-78*

Batch execution queues

- See also Execution queues
- description, *Vol 1, 14-3*

Batch identifiers, *Vol 1, 12-11*

Batch jobs

- See also Batch processing environment; Batch queues
- accessing devices, *Vol 1, 9-61*
- allowing to complete before stopping a queue, *Vol 1, 14-63*
- changing scheduling priority, *Vol 1, 14-36*
- completing before using VMSINSTAL.COM command procedure, *Vol 1, 3-3*
- controlling, *Vol 1, 14-65*
- deleting, *Vol 1, 14-69*
- distributing system work load, *Vol 2, 17-3*
- executing, *Vol 1, 14-3*
- holding and releasing, *Vol 1, 14-67*
- job card, *Vol 1, 8-39*
- modifying, *Vol 1, 14-67*
- monitoring, *Vol 1, 14-65*
- requeuing
 - executing, *Vol 1, 14-68*
 - pending, *Vol 1, 14-69*
- retaining in a queue, *Vol 1, 14-29*
- scheduling, *Vol 1, 14-36*
- status
 - See Job status
- submitting at startup, *Vol 1, 5-14*

Batch mode

- as execution mode for a startup procedure, *Vol 1, 5-20*

Batch processing environment

- See also Batch jobs; Batch queues
- generic queues in an OpenVMS Cluster environment, *Vol 1, 14-6*
- on a standalone workstation, *Vol 1, 14-5*
- sample configurations, *Vol 1, 14-5 to 14-7*
- steps for setting up, *Vol 1, 14-13*
- with specialized queues, *Vol 1, 14-5*

Batch queues

- See also Batch jobs; Batch processing environment
- allowing jobs to complete before stopping, *Vol 1, 14-63*
- assigning to resource affinity domain (RAD), *Vol 1, 14-35*
- commands for managing, *Vol 1, 14-57*
- creating, *Vol 1, 14-15*
- deleting, *Vol 1, 14-65*
- for memory constrained systems, *Vol 1, 14-35*

Batch queues (cont'd)

- on standalone workstations, *Vol 1, 14-5*
- optimizing for Sort/Merge utility, *Vol 1, 14-35*
- options, *Vol 1, 14-20*
 - controlling job performance and resources, *Vol 1, 14-32*
 - qualifiers for specifying, *Vol 1, 14-20 to 14-21*
 - restricting access, *Vol 1, 14-23*
 - retaining jobs, *Vol 1, 14-27*
- pausing, *Vol 1, 14-61*
- setting up for vector processing, *Vol 2, 28-7*
- starting, *Vol 1, 14-15*
- status, *Vol 1, 14-59*
- stopping, *Vol 1, 14-62, 14-63*
 - before shutting down a node, *Vol 1, 14-63*

Battery-Backed Watch

- setting time using, *Vol 1, 6-17*

Binary Error Log Translation utility

- in DECEvent kit, *Vol 2, 20-2*

Binding volumes into a volume set, *Vol 1, 9-29*

BIOLM process limit, *Vol 1, 7-44*

- value for efficient backups, *Vol 1, 11-11*

BITMAP.SYS file

- See Storage bitmap file

Bitmaps

- index file, *Vol 1, 9-5*
- storage, *Vol 1, 9-5*

Blocks

- disk
 - definition, *Vol 1, 9-3*
 - erasing, *Vol 1, 9-73*
- Files-11
 - definition, *Vol 2, A-1*
- tape
 - definition, *Vol 1, 9-11*

Boot block

- definition, *Vol 1, 4-19*
- in index file, *Vol 2, A-6*
- processors using, *Vol 1, 4-19*
- role in boot process, *Vol 1, 4-2*
- writing with Writeboot utility, *Vol 1, 4-19*

Booting

- automatically after shutting down, *Vol 1, 4-24*
- bootstrap image
 - Alpha, *Vol 1, 4-21*
 - in index file, *Vol 2, A-6*
 - VAX, *Vol 1, 4-20*
- conversationally
 - See Conversational boot
- cross-architecture, *Vol 2, 24-28*
- definition, *Vol 1, 4-2*
- displaying startup commands, *Vol 1, 4-16*
- from an alternate system disk, *Vol 1, 4-4*
- in a multiprocessing system, *Vol 2, 28-2*
- in an emergency
 - with default system parameters, *Vol 1, 4-9*

Booting
 in an emergency (cont'd)
 without startup and login procedures, *Vol 1*, 4–10
 without the user authorization file, *Vol 1*, 4–11
 installation of page and swap files, *Vol 2*, 16–6
 loading VVIEF code, *Vol 2*, 28–5
 location of computer-specific instructions, *Vol 1*, 4–2
 message
 See Boot messages
 nonstop, *Vol 1*, 4–4
 problems
 fixing by booting with default parameter values, *Vol 1*, 4–9
 fixing by booting with minimum startup, *Vol 1*, 4–15
 hardware, *Vol 1*, 4–18
 invalid boot block, *Vol 1*, 4–19
 solving, *Vol 1*, 4–18
 use of boot block, *Vol 1*, 4–19
 with an alternate system parameter file, *Vol 1*, 4–8
 with controlled system startup, *Vol 1*, 4–13
 Boot messages
 indicating execution of STARTUP.COM procedure, *Vol 1*, 4–6
 indicating execution of STARTUP_VMS.COM procedure, *Vol 1*, 4–6
 indicating success, *Vol 1*, 4–6
 indicating that login is possible, *Vol 1*, 4–6
 question mark (?), *Vol 1*, 4–18
 Bootstrapping
 See Booting
 BOT markers, *Vol 1*, 9–11
 Break-ins
 auditing attempts, *Vol 2*, 20–27
 BROADCAST device setting, *Vol 2*, 20–25
 Broadcast messages
 removing, *Vol 2*, 22–10
 Buffered I/O
 byte count limit, *Vol 1*, 7–44
 count limit, *Vol 1*, 7–44
 Buffers
 using large buffers in LAT environment, *Vol 2*, 26–13
 Bugcheck
 message during UETP, *Vol 2*, 19–31
 Burst bars, *Vol 1*, 14–39
 Burst pages, *Vol 1*, 14–39
 file, *Vol 1*, 14–43
 job, *Vol 1*, 14–42
 Busy queue status, *Vol 1*, 14–59
 BYTLM process limit
 value for efficient backups, *Vol 1*, 11–11

BYTLM quota, *Vol 1*, 7–44

C

C2-class system
 installing OpenVMS to use, *Vol 1*, 3–2
 Caches
 Extended File Cache, *Vol 2*, 18–1
 virtual I/O, *Vol 2*, 18–10
 Virtual I/O Cache, *Vol 2*, 18–1
 CACHE_SERVER process
 creation during system startup, *Vol 1*, 5–5
 Caching
 ACP system parameters, *Vol 1*, 9–57
 file system data, *Vol 2*, 18–1
 managing file system data caches, *Vol 2*, 18–1
 Caching attribute
 defined, *Vol 2*, 18–7
 setting, *Vol 2*, 18–7
 showing, *Vol 2*, 18–7
 Canceling
 characteristics on a queue, *Vol 1*, 14–31
 Capability of a vector
 See Vector capability
 Captive accounts, *Vol 1*, 7–31
 Card readers
 operating, *Vol 1*, 8–39
 translation modes, *Vol 1*, 8–40
 Cards
 decks, *Vol 1*, 8–39
 CD-R (CD-Recordable) drive, *Vol 1*, 10–41
 CDRECORD.COM, *Vol 1*, 10–41
 CD-ROM (compact disc read-only memory), *Vol 1*, 9–3 to 9–5
 accessing ISO 9660-formatted, *Vol 1*, 9–47
 accessing with the InfoServer Client for OpenVMS, *Vol 2*, 25–11
 characteristics of, *Vol 1*, 9–2
 file structures, *Vol 1*, 9–3
 formats for automatic serving, *Vol 2*, 25–2
 High Sierra format, *Vol 1*, 9–3
 ISO 9660 format, *Vol 1*, 9–3
 media formats used, *Vol 1*, 9–4
 CD-ROM (compact disc read-only memory)
 creating your own, *Vol 1*, 10–41
 CD-RW (CD-Record/Write) drive, *Vol 1*, 10–41
 Changing Compaq supplied data in the Help Message database, *Vol 1*, 5–31
 Changing system parameters
 recommended method, *Vol 2*, 15–4, 15–18
 with AUTOGEN, *Vol 2*, 15–18
 with conversational boot, *Vol 1*, 4–7
 with SYSMAN, *Vol 2*, 15–29
 CIPCA adapter
 system tuning, *Vol 2*, 15–37

- CIXCD adapter
 - system tuning, *Vol 2*, 15–37
- Classes
 - enabling and disabling, *Vol 2*, 20–24
- Class of data
 - in SHOW CLUSTER, *Vol 2*, 22–5
 - removing, *Vol 2*, 22–11
- Class scheduler, *Vol 1*, 7–29
- Closed queue status, *Vol 1*, 14–59
- Closing a queue, *Vol 1*, 14–62
- Cluster failover of OpenVMS Registry server, *Vol 2*, D–9
- Cluster MIB
 - defined, *Vol 2*, C–11
 - setting up, *Vol 2*, C–13
 - subagent objects, *Vol 2*, C–11
 - subagents, *Vol 2*, C–11
- Clusters
 - See OpenVMS Cluster systems
- CLUSTER_CONFIG.COM command procedure, *Vol 1*, 5–6; *Vol 2*, 22–2
 - compared to VMSKITBLD.COM command procedure, *Vol 1*, 2–32
 - creating SATELLITE_PAGE.COM command procedure, *Vol 2*, 16–26
 - setting up LAN MOP, *Vol 2*, 24–26
 - use in adding a system to an OpenVMS Cluster environment, *Vol 1*, 2–32
- CLUSTER_SERVER process
 - creation during system startup, *Vol 1*, 5–4
- CLUSTER_SIZE attribute, *Vol 1*, 11–61
- Code compilers, *Vol 1*, 10–13
- Command files
 - LAN Control Program (LANCP) utility, *Vol 2*, 24–10
- Command formats
 - for backups, *Vol 1*, 11–5
 - for image backups, *Vol 1*, 11–33, 11–34
 - for incremental backups, *Vol 1*, 11–36, 11–37
 - for multiple backup output devices, *Vol 1*, 11–26, 11–34, 11–35
- Command procedures
 - executing in SYSMAN, *Vol 1*, 2–18, 2–19
 - for backups, *Vol 1*, 11–39
 - for image backups, *Vol 1*, 11–39
 - for incremental backups, *Vol 1*, 11–41
 - for installing products, *Vol 1*, 3–14, 3–29
 - See also VMSINSTAL.COM command procedure
 - for interactive backups, *Vol 1*, 11–42
 - for MONITOR, *Vol 2*, 20–40 to 20–43
 - for setting up storage media, *Vol 1*, 9–60
 - disk volumes, *Vol 1*, 9–60
 - tape volumes, *Vol 1*, 9–61
 - for SHOW CLUSTER, *Vol 2*, 22–15
 - controlling output, *Vol 2*, 22–6
 - default file type, *Vol 2*, 22–16
- Command procedures
 - for SHOW CLUSTER (cont'd)
 - description, *Vol 2*, 22–16
 - formatting reports, *Vol 2*, 22–10
 - initialization, *Vol 2*, 22–14
 - SHOW_CLUSTER\$INIT.COM, *Vol 2*, 22–14
 - for system management, *Vol 1*, 2–5
 - for system startup, *Vol 1*, 2–12
 - See also Startup command procedure LIBDECOMP.COM, *Vol 2*, 17–9
 - login, *Vol 1*, 2–17
 - setting protection in, *Vol 1*, 10–22
 - testing a spooled printer, *Vol 1*, 8–37
 - to run AUTOGEN regularly, *Vol 2*, 15–23
 - to specify default state of operator log files, *Vol 2*, 20–24
 - VVIEF\$INSTAL.COM procedure, *Vol 2*, 28–5
- Commands
 - See DCL commands
- Communications
 - with operators, *Vol 1*, 9–25
 - with users, *Vol 1*, 9–25
 - creating systemwide announcements, *Vol 1*, 5–14
- Compact disc drives
 - supported by UETP, *Vol 2*, 19–7
- Compact disc read-only memory
 - See CD-ROM
- Compaq support representative
 - reporting system failure, *Vol 2*, 16–19
- Compaq TCP/IP Services for OpenVMS
 - description, *Vol 2*, 23–2
 - installing and configuring, *Vol 2*, 23–6
 - management tools, *Vol 2*, 23–5
 - preparing to join a network, *Vol 2*, 23–6
 - related documentation, *Vol 2*, 23–7
 - starting and stopping, *Vol 2*, 23–7
 - summary of benefits, *Vol 2*, 23–4
 - support for clusters, *Vol 2*, 23–5
- Compare operations
 - with BACKUP, *Vol 1*, 11–27
- Completion status
 - showing for batch and print jobs, *Vol 1*, 14–27
- Compressing the Help Message database after deletions, *Vol 1*, 5–30
- Configurations
 - displaying LAN, *Vol 2*, 24–10
 - for software product options, *Vol 1*, 3–19
 - LAN, *Vol 2*, 24–1
 - queue
 - sample batch queuing system, *Vol 1*, 14–5 to 14–7
 - sample print queuing system, *Vol 1*, 14–8 to 14–13

CONFIGURATION SET command
 in SYSMAN, *Vol 2*, 22–16
 CONFIGURATION SHOW command
 in SYSMAN, *Vol 2*, 22–16
 CONFIGURE phase of system startup, *Vol 1*, 5–4,
 5–18, 8–10
 definition, *Vol 1*, 8–7
 CONFIGURE process
 starting, *Vol 1*, 5–5
 Configuring devices
 HSC
 disabling during system startup, *Vol 1*,
 8–10
 in system startup, *Vol 1*, 5–7, 8–7
 CONINTERR.EXE driver, *Vol 1*, 8–9
 CONNECT command
 See also IO CONNECT command
 in SYSGEN (VAX), *Vol 1*, 8–8
 for connecting the console device, *Vol 1*,
 8–8
 in system startup, *Vol 1*, 5–7
 Connecting devices
 automatically, *Vol 1*, 8–7, 8–9
 in system startup, *Vol 1*, 5–4
 manually, *Vol 1*, 8–9
 in system startup, *Vol 1*, 5–7
 on VAX, *Vol 1*, 8–8
 network communications device, *Vol 1*, 8–10
 on VAX systems, *Vol 1*, 8–9
 virtual terminals, *Vol 1*, 8–22
 Connecting to a Windows NT system, *Vol 2*, D–9
 CONSCOPY.COM command procedure, *Vol 1*,
 5–35
 Console report during UETP, *Vol 2*, 19–15
 Console storage device
 connecting (VAX), *Vol 1*, 8–8
 copying, *Vol 1*, 5–35
 use in booting, *Vol 1*, 4–2
 Console terminals, *Vol 1*, 2–22, 2–24
 message
 indicating lack of installed page file, *Vol 1*,
 5–6
 login welcome, *Vol 1*, 2–10
 Container file
 for creating a CD-ROM, *Vol 1*, 10–42
 reusing, *Vol 1*, 10–44
 CONTIN.SYS file
 See Continuation file
 Continuation file
 CONTIN.SYS, *Vol 2*, A–9
 reserved file, *Vol 2*, A–9
 used as extension file identifier, *Vol 2*, A–9
 Continuation volumes
 in volume set, *Vol 1*, 9–54
 mounting in tape volume sets, *Vol 1*, 9–52
 CONTINUE command
 in conversational boot, *Vol 1*, 4–7
 Control access
 for disk files, *Vol 1*, 10–21, 10–26
 Controller configuration problems, *Vol 2*, 25–12
 Conversational boot
 changing system parameters, *Vol 1*, 4–7; *Vol*
 2, 15–36
 CONTINUE command, *Vol 1*, 4–7
 location of computer-specific instructions, *Vol 1*,
 4–4
 performing, *Vol 1*, 4–4
 SET/STARTUP command, *Vol 1*, 4–14
 SET command, *Vol 1*, 4–7
 SHOW/STARTUP command, *Vol 1*, 4–14
 SHOW command, *Vol 1*, 4–7
 showing system parameters, *Vol 1*, 4–7; *Vol 2*,
 15–36
 specifying an alternate startup command
 procedure, *Vol 1*, 4–14
 SYSBOOT prompt, *Vol 1*, 4–4
 tasks allowed in, *Vol 1*, 4–4
 USE command, *Vol 1*, 4–8
 uses of, *Vol 1*, 4–4
 with alternate system parameter file, *Vol 1*,
 4–8
 with minimum startup, *Vol 1*, 4–15
 Convert utility (CONVERT)
 saving the queue database, *Vol 1*, 13–13
 using to change organization of file, *Vol 1*,
 10–34
 Coordinated Universal Time (UTC), *Vol 1*, 6–1
 COPY command, *Vol 1*, 11–24
 disk volumes, *Vol 1*, 10–35
 in System Dump Analyzer utility, *Vol 2*, 16–5,
 16–22, 16–23
 restriction for copying dump files, *Vol 2*, 16–22
 sending message after file is copied, *Vol 1*,
 10–38
 standard-labeled volumes
 copying from, *Vol 1*, 10–35
 tape volumes
 copying files, *Vol 1*, 10–35, 10–37
 transferring information, *Vol 1*, 10–34
 Copying directories
 with BACKUP, *Vol 1*, 11–24
 Copying files
 from disk volumes, *Vol 1*, 10–35
 from tape volumes, *Vol 1*, 10–35
 methods for, *Vol 1*, 10–34
 system dump files, *Vol 2*, 16–23
 to disk volumes, *Vol 1*, 10–35
 to tape volumes, *Vol 1*, 10–37
 using COPY command, *Vol 1*, 10–34
 using Exchange utility, *Vol 1*, 10–40
 with BACKUP, *Vol 1*, 11–24

- Copying software product release notes, *Vol 1*, 3–25
- Core image file
 - CORIMG.SYS, *Vol 2*, A–9
 - not supported by OpenVMS, *Vol 2*, A–9
- CORIMG.SYS file
 - See Core image file
- Counters
 - status of LAT node, *Vol 2*, 26–15
- CPUDEFAULT process limit
 - choosing a value for batch queues, *Vol 1*, 14–34
 - specifying a value for batch queues, *Vol 1*, 14–20, 14–32
- CPU identification number, *Vol 2*, 28–2
- CPUMAXIMUM process limit
 - choosing a value for batch queues, *Vol 1*, 14–34
 - specifying a value for batch queues, *Vol 1*, 14–21, 14–32
- CPU process limit, *Vol 1*, 7–44
- CPU time
 - restricting, *Vol 1*, 7–29
- CRASH console command, *Vol 1*, 4–22
- Crash dumps, *Vol 2*, 16–2
 - See also System dump files; System failures analyzing, *Vol 2*, 16–19
 - analyzing in system startup, *Vol 1*, 5–13
 - comparison of physical and selective, *Vol 2*, 16–4
 - freeing page file of, *Vol 2*, 16–24
 - in system startup, *Vol 2*, 16–22
 - physical, *Vol 2*, 16–4
 - releasing, *Vol 2*, 16–24, 16–25
 - requirements for saving, *Vol 2*, 16–4
 - saving contents of page file on reboot, *Vol 2*, 16–3
 - saving contents of system dump file on reboot, *Vol 1*, 5–13
 - selective, *Vol 2*, 16–4
- Crash Log Utility Extractor (CLUE)
 - description, *Vol 2*, 16–21
- CREATE/DIRECTORY command
 - to specify UIC-based directory protection, *Vol 1*, 10–27
- CREATE command
 - creating directories, *Vol 1*, 10–35
 - limiting number of file versions, *Vol 1*, 9–68
 - in SYSGEN
 - changing page, swap, and dump file sizes, *Vol 2*, 16–33
 - creating page, swap, and dump files, *Vol 2*, 16–33
 - writing new file to tape volume, *Vol 1*, 10–33
- Creating
 - additional queue managers, *Vol 1*, 13–10
 - a PCF during installation, *Vol 1*, 3–26
 - execution queues, *Vol 1*, 14–15
- Creating
 - execution queues (cont'd)
 - autostart, *Vol 1*, 14–15, 14–17
 - nonautostart, *Vol 1*, 14–18
 - forms, *Vol 1*, 14–49
 - generic queues, *Vol 1*, 14–19
 - operator log files, *Vol 2*, 20–23
 - page, swap, and dump files
 - with AUTOGEN, *Vol 2*, 16–31
 - with SYSGEN, *Vol 2*, 16–33
 - queue databases, *Vol 1*, 13–7
 - system parameter files
 - with SYSGEN, *Vol 2*, 15–35
- Creating search path of Help Message database files, *Vol 1*, 5–28
- Cross-architecture
 - booting, *Vol 2*, 24–28
- CTRLNAME logical name, *Vol 2*, 19–33
- Current accounting file
 - controlling which resources are tracked in, *Vol 2*, 21–3
 - default file name, *Vol 2*, 21–3
 - definition, *Vol 2*, 21–1
 - finding out what resources are tracked in, *Vol 2*, 21–2
 - moving, *Vol 2*, 21–3
 - starting up a new version of, *Vol 2*, 21–3
- Current system parameters, *Vol 2*, 15–3, 15–26
- Customizing the system
 - adding optional system files that have been removed from the system disk, *Vol 1*, 5–1
 - backing up
 - console storage device, *Vol 1*, 5–35
 - system disk, *Vol 1*, 5–35
 - building standalone BACKUP, *Vol 1*, 5–35
 - creating site-specific startup command procedures, *Vol 1*, 5–2
 - creating systemwide announcements, *Vol 1*, 5–14
 - duplicating the system disk, *Vol 1*, 2–27
 - enabling autostart, *Vol 1*, 5–12
 - installing known images, *Vol 1*, 5–13
 - installing resident images (Alpha), *Vol 1*, 5–13
 - limiting the number of interactive users, *Vol 1*, 5–16
 - making remote InfoServer devices available, *Vol 2*, 25–16
 - making remote InfoServer disks available, *Vol 1*, 5–13
 - modifying login procedures, *Vol 1*, 5–17
 - modifying site-specific startup command procedures, *Vol 1*, 5–2
 - rules, *Vol 1*, 5–4
 - SYCONFIG.COM, *Vol 1*, 5–7
 - SYLOGICALS.COM, *Vol 1*, 5–8
 - SYPAGSWPFILES.COM, *Vol 1*, 5–6
 - SYSECURITY.COM, *Vol 1*, 5–10
 - SYSTARTUP_VMS.COM, *Vol 1*, 5–10

Customizing the system (cont'd)

- removing optional system files from the system disk, *Vol 1*, 5-1
- setting up a LAT network, *Vol 1*, 5-15
- starting InfoServer Client for OpenVMS, *Vol 1*, 5-13; *Vol 2*, 25-11
- starting queues, *Vol 1*, 5-12
- starting the DECnet network, *Vol 1*, 5-16
- submitting batch jobs at system startup, *Vol 1*, 5-14

Cylinders

- definition, *Vol 2*, A-2

D

DAD virtual disk unit, *Vol 2*, 25-14

Databases

- LAT database, *Vol 2*, 26-14
- LMF
 - use in system startup, *Vol 1*, 5-5
- of software products, *Vol 1*, 3-16
- queue
 - See Queue database
- startup
 - definition, *Vol 1*, 5-18
 - layered product, *Vol 1*, 5-19
 - OpenVMS, *Vol 1*, 5-18
 - order of startup events, *Vol 1*, 5-4

Data blocks

- partially recorded
 - ISO 9660 standard, *Vol 1*, 9-8

Data card deck, *Vol 1*, 8-40

Data Communications Equipment, *Vol 1*, 8-27, 8-28

Data interleaving

- ISO 9660, *Vol 1*, 9-8

Data loss

- avoiding by dismounting volume, *Vol 1*, 9-58

Data modified, *Vol 1*, 10-18

Date formats, *Vol 1*, 6-17

- predefined, *Vol 1*, 6-21
- specifying, *Vol 1*, 6-19, 6-20

Daylight saving time, *Vol 1*, 6-2

- changing to, *Vol 1*, 6-13
- automatically, *Vol 1*, 6-13

DBBFs (detected bad block files), *Vol 1*, 9-80

DCL

- access dates, *Vol 1*, 10-18

DCL commands

- accessing disk and tape files, *Vol 1*, 10-28
- executing on multiple nodes, *Vol 1*, 2-18
- executing with SYSMAN DO command, *Vol 1*, 2-18
- file manipulation, *Vol 1*, 10-16
- for system management, *Vol 1*, 2-5
- retrieving file information, *Vol 1*, 10-16
- with DO command in SYSMAN, *Vol 2*, 22-17

DCL interface, *Vol 1*, 3-14

DCL statements

- filtering from report, *Vol 2*, 15-12

DCOM, *Vol 1*, 10-9

ddcu format

- for device names, *Vol 1*, 8-1

DEALLOCATE command, *Vol 1*, 9-16

- tape, *Vol 1*, 10-38

Deallocating devices, *Vol 1*, 9-16

Deallocating drives, *Vol 1*, 9-16

DECdns (Digital Distributed Name Service)

- name service
 - storing full names in, *Vol 2*, 23-9

DECdtm services

- and managing transaction logs, *Vol 2*, 27-1
- disabling, *Vol 2*, 27-23
- enabling, *Vol 2*, 27-23
- using in a DECnet-Plus network, *Vol 2*, 27-4

DECdtm transaction logs

- changing the size of, *Vol 2*, 27-12
- checking the size of, *Vol 2*, 27-11
- creating, *Vol 2*, 27-6
- definition, *Vol 2*, 27-3
- moving, *Vol 2*, 27-14
- planning, *Vol 2*, 27-3

DECEvent

- Binary Error Log Translation utility, *Vol 2*, 20-2

DECEvent utility

- additional commands, *Vol 2*, 20-12
- description, *Vol 2*, 20-3
- DIAGNOSE command, *Vol 2*, 20-3
- introduction, *Vol 2*, 20-8
- invoking, *Vol 2*, 20-10
- privileges required, *Vol 2*, 20-12
- producing reports, *Vol 2*, 20-12
- qualifiers, *Vol 2*, 20-10
- report formats, *Vol 2*, 20-9
 - brief, *Vol 2*, 20-13
 - fast error (FSTERR), *Vol 2*, 20-16
 - full, *Vol 2*, 20-12
 - summary, *Vol 2*, 20-15
 - terse, *Vol 2*, 20-14
- using to report on error log file, *Vol 2*, 20-9

DECnet

- asynchronous
 - using virtual terminals, *Vol 1*, 8-22, 8-23
- error message during UETP, *Vol 2*, 19-29
- managing remote nodes with, *Vol 1*, 2-13
- overriding AUTOGEN observations, *Vol 2*, 15-21
- over TCP/IP, *Vol 2*, 23-3
 - enabling, *Vol 2*, 23-12
- preparing for UETP, *Vol 2*, 19-12
- specifying transport, *Vol 1*, 5-34
- starting, *Vol 1*, 5-16
- UETP defaults for installation, *Vol 2*, 19-31
- UETP test phase, *Vol 2*, 19-36, 19-37

DECnet for OpenVMS
 comparing with DECnet-Plus for OpenVMS, *Vol 2*, 23–9
 description, *Vol 2*, 23–2, 23–3
 node names, *Vol 2*, 23–9
 remote management of node, *Vol 2*, 23–10

DECnet Phase IV
 See DECnet for OpenVMS

DECnet Phase V
 See DECnet-Plus for OpenVMS

DECnet-Plus for OpenVMS
 advantages, *Vol 2*, 23–9
 comparing with DECnet for OpenVMS, *Vol 2*, 23–9
 description, *Vol 2*, 23–2
 distinguished from DECnet for OpenVMS, *Vol 2*, 23–2
 features, *Vol 2*, 23–9
 full names, *Vol 2*, 23–9
 installing and configuring, *Vol 2*, 23–11
 management tools, *Vol 2*, 23–10
 moving your network to, *Vol 2*, 23–13
 new features, *Vol 2*, 23–9
 preparing to join a network, *Vol 2*, 23–11
 related documentation, *Vol 2*, 23–13
 starting, *Vol 1*, 5–16
 starting and stopping, *Vol 2*, 23–13
 summary of benefits, *Vol 2*, 23–8
 support for clusters, *Vol 2*, 23–10
 transition
 definition, *Vol 2*, 23–13
 using DECnet over TCP/IP, *Vol 2*, 23–12

Decompressing system libraries
 See Expanding system libraries

DECW\$TAILOR
 See Tailoring utilities

Deductible resource, *Vol 1*, 7–3

DEFAULT
 account in UAF, *Vol 1*, 7–6
 Default boot procedure, *Vol 1*, 4–2
 Default directories, *Vol 1*, 7–15
 Default form, *Vol 1*, 14–50
 DEFAULT form, *Vol 1*, 14–48
 Default protection
 UIC, *Vol 1*, 10–22
 Default Protection ACEs, *Vol 1*, 12–10
 Default system parameters
 booting with, *Vol 1*, 4–9
 Deferred memory testing
 on AlphaServer 4100 computers, *Vol 1*, 4–3

DEFINE/CHARACTERISTIC command, *Vol 1*, 14–30

DEFINE/FORM command, *Vol 1*, 14–49
 for controlling line overflow, *Vol 1*, 14–52

Defragmenting disks, *Vol 1*, 11–51

Deinstalling layered software products, *Vol 1*, 3–31

DELETE/CHARACTERISTIC command, *Vol 1*, 14–32

DELETE/ENTRY command, *Vol 1*, 14–75

DELETE/FORM command, *Vol 1*, 14–51

DELETE/QUEUE command, *Vol 1*, 14–65

Delete access
 explicitly assigning, *Vol 1*, 10–26
 for disk directory files, *Vol 1*, 10–26
 for disk files, *Vol 1*, 10–21
 granting through protection codes, *Vol 1*, 12–8

Deleting
 Compaq messages from the Help Message database, *Vol 1*, 5–29
 files
 from the system disk, *Vol 1*, 5–1
 with PRINT/DELETE command, *Vol 1*, 14–78
 forms, *Vol 1*, 14–51
 problems with, *Vol 1*, 14–76
 jobs, *Vol 1*, 14–69
 page, swap, and dump files
 after creating new version, *Vol 2*, 16–35
 caution, *Vol 2*, 16–28
 queue characteristics, *Vol 1*, 14–32
 problems with, *Vol 1*, 14–76
 queues, *Vol 1*, 14–65
 problems with, *Vol 1*, 14–76

Delta/XDelta Debugger (DELTA/XDELTA)
 entering through IPC, *Vol 1*, 9–80

Dependent software products, *Vol 1*, 3–24

DESELECT command
 in SHOW CLUSTER, *Vol 2*, 22–13

Despooling a spooled printer, *Vol 1*, 8–37

Destination for installing software
 copying files, *Vol 1*, 3–30

Destination parameter
 in VMSINSTAL.COM, *Vol 1*, 3–9

Detected bad block files (DBBFs)
 See DBBFs

Device
 definition, *Vol 1*, 9–2

Device control libraries, *Vol 1*, 14–52 to 14–54
 See also Device control modules
 assigning to a queue, *Vol 1*, 14–54
 order of module output, *Vol 1*, 14–53
 sample commands, *Vol 1*, 14–56
 setting up, *Vol 1*, 14–52

Device control modules, *Vol 1*, 14–52 to 14–79
 See also Device control libraries
 adding, *Vol 1*, 14–54, 14–79
 creating, *Vol 1*, 14–54
 deleting, *Vol 1*, 14–54, 14–79
 forms, *Vol 1*, 14–55
 inserting into a library, *Vol 1*, 14–54
 listing, *Vol 1*, 14–55

Device control modules (cont'd)

- naming, *Vol 1, 14-55*
- order of output, *Vol 1, 14-53*
- page setup, *Vol 1, 14-52, 14-55*
- requesting with PRINT command, *Vol 1, 14-55*
- reset, *Vol 1, 14-55*
 - when queue is started, *Vol 1, 14-56*
- sample commands, *Vol 1, 14-56*
- setting up, *Vol 1, 14-53*
- setup, *Vol 1, 14-52, 14-55*
- specifying, *Vol 1, 14-53*
- storing, *Vol 1, 14-55*
- troubleshooting, *Vol 1, 14-79*
- types, *Vol 1, 14-52*
- with forms, *Vol 1, 14-53*

Device drivers

- CONINTERR.EXE file, *Vol 1, 8-9*
- for event handling, *Vol 1, 8-9*
- loading
 - automatically, *Vol 1, 8-7*
 - in system startup, *Vol 1, 5-4, 5-7*
 - manually (Alpha), *Vol 1, 8-9*
 - manually (VAX), *Vol 1, 8-8*
- not associated with a specific device, *Vol 1, 8-9*
- TTDRIVER, *Vol 1, 8-22*

Device names, *Vol 1, 8-1*

- for virtual terminals, *Vol 1, 8-22*
- in an OpenVMS Cluster environment, *Vol 1, 8-1*

DEVICE phase of system startup, *Vol 1, 5-4, 5-18*

Devices

- accessing in batch job, *Vol 1, 9-61*
- allocating, *Vol 1, 9-15, 10-38*
- autoconfiguration of
 - device descriptions, *Vol 1, 8-13*
 - supported buses, *Vol 1, 8-17*
- availability
 - OPCOM message, *Vol 1, 9-73*
- configuring
 - automatically, *Vol 1, 8-11*
 - in system startup, *Vol 1, 5-7, 8-7*
 - manually, *Vol 1, 5-7, 8-8*
 - special devices, *Vol 1, 8-8, 8-9*
- connecting, *Vol 1, 8-8*
- determining available, *Vol 1, 8-2*
- Ethernet adapter
 - specifying number for AUTOGEN, *Vol 2, 15-21*
- getting information about, *Vol 1, 8-2, 8-38*
- ISO-9660
 - getting information about, *Vol 1, 8-5*
- LTA_n, *Vol 1, 8-23*
- magnetic tape
 - See Tapes
- managing, *Vol 1, 8-1*
- managing with LAN Control Program (LANCP) utility, *Vol 2, 24-10*

Devices (cont'd)

- manually configuring non-standard devices, *Vol 1, 5-7, 8-8*
 - mounting volumes, *Vol 1, 9-25*
 - network communications
 - connecting, *Vol 1, 8-10*
 - connecting (VAX), *Vol 1, 8-9*
 - not recognized by the system, *Vol 1, 8-8*
 - OPA0:, *Vol 1, 2-24*
 - printers
 - See Printers
 - protecting, *Vol 1, 8-6*
 - requiring manual connecting, *Vol 1, 8-10*
 - resetting default in SYSMAN, *Vol 1, 2-18*
 - RTA_n, *Vol 1, 8-23*
 - setting characteristics, *Vol 1, 14-14*
 - in system startup, *Vol 1, 14-14*
 - special
 - connecting, *Vol 1, 8-9*
 - spooled, *Vol 1, 14-14*
 - status report on, *Vol 2, 20-18*
 - suppressing autoconfiguration during system startup, *Vol 1, 5-8, 8-10*
 - terminals
 - See Terminals
- ## Device tests
- running individually with UETP, *Vol 2, 19-32*
- ## Device unavailable queue status, *Vol 1, 14-59*
- ## DIAGNOSE command, *Vol 2, 20-3*
- ## Diagnostics
- relationship to UETP, *Vol 2, 19-16*
- ## Dialup identifiers, *Vol 1, 12-11*
- ## DIBOL
- starting the message manager, *Vol 1, 5-16*
- ## DIGITAL network architecture
- See DNA
- ## DIGITAL System Identifier
- See DSI
- ## DIOLM process limit, *Vol 1, 7-44*
- value for efficient backups, *Vol 1, 11-11*
- ## Direct I/O count process limit, *Vol 1, 7-44*
- ## Direct mode
- as execution mode for a startup procedure, *Vol 1, 5-20*
- ## Directories
- backing up, *Vol 1, 11-26*
 - backlink, *Vol 1, 9-71*
 - copying with BACKUP, *Vol 1, 11-24*
 - creating, *Vol 1, 10-35*
 - destination
 - specifying in VMSINSTAL.COM command procedure, *Vol 1, 3-9*
 - for an interactive account, *Vol 1, 7-15*
 - levels of access in restore operations, *Vol 1, 11-31*
 - protecting, *Vol 1, 10-26, 10-27*
 - resetting default in SYSMAN, *Vol 1, 2-18*

Directories (cont'd)

- working
 - in VMSINSTAL.COM command procedure, *Vol 1*, 3–10
- DIRECTORY command, *Vol 1*, 10–20
 - checking number of user's disk blocks, *Vol 1*, 9–63
 - retrieving file information, *Vol 1*, 10–16
 - showing full information, *Vol 1*, 10–32
 - to obtain product list, *Vol 1*, 3–7
 - with tapes, *Vol 1*, 10–31, 10–35
- Directory trees
 - copying, *Vol 1*, 11–24
- DISABLE AUTOSTART/QUEUES command, *Vol 1*, 14–62, 14–63
 - entering before shutting down a system, *Vol 1*, 14–63
 - relationship to STOP/QUEUES/ON_NODE command, *Vol 1*, 14–63
- Disabled queues, *Vol 1*, 14–80
- Disabling autostart, *Vol 1*, 13–9, 14–62
 - before shutting down a node, *Vol 1*, 14–63
- Disabling caching
 - clusterwide, *Vol 2*, 18–3
 - for a volume, *Vol 2*, 18–3
 - read-ahead, *Vol 2*, 18–8
- Disabling operator terminals, *Vol 1*, 2–25
- Disabling spooling, *Vol 1*, 8–37
- Disabling user accounts, *Vol 1*, 7–27
- Disk defragmenters, *Vol 1*, 10–13
- Disk files
 - accessing at file level, *Vol 1*, 10–28
 - assigning an alias, *Vol 1*, 10–25
 - copying
 - from disk volumes, *Vol 1*, 10–35
 - to tapes, *Vol 1*, 10–37
 - using COPY command, *Vol 1*, 10–34
 - modifying characteristics, *Vol 1*, 10–24, 10–25
- Disk I/O resource
 - offloading
 - virtual I/O cache, *Vol 2*, 18–10
- DISKQUOTA/DISABLE command, *Vol 1*, 9–67
- DISKQUOTA/ENABLE command, *Vol 1*, 9–67
- DISKQUOTA commands
 - See also Disk Quota utility
 - in SYSMAN, *Vol 1*, 2–12
- Disk quotas, *Vol 1*, 7–15
 - See also Disk Quota utility
 - creating, *Vol 1*, 9–65
 - definition, *Vol 1*, 9–63
 - disabling, *Vol 1*, 9–67
 - displaying, *Vol 1*, 9–63
 - ensuring accuracy with rebuild operation, *Vol 1*, 9–64
 - establishing, *Vol 1*, 9–65
 - example, *Vol 1*, 7–36
 - exceeding, *Vol 1*, 9–64
 - file, *Vol 1*, 9–63

Disk quotas (cont'd)

- maintaining, *Vol 1*, 9–64
 - retrieving information, *Vol 1*, 9–66
 - suspending operations, *Vol 1*, 9–67
- Disk Quota utility (DISKQUOTA), *Vol 1*, 9–65
- ## Disks
- See also CD-ROM; Disk files; Disk quotas; Disk volumes; System disks
- allocating drives, *Vol 1*, 9–15
 - allocating space on, *Vol 1*, 9–66
 - backing up active, *Vol 1*, 11–66
 - block
 - definition, *Vol 1*, 9–3
 - grouped into cluster, *Vol 2*, A–1
 - changing volume label, *Vol 1*, 3–30
 - cluster
 - definition, *Vol 1*, 9–3
 - unit of allocation, *Vol 2*, A–1
 - concepts, *Vol 2*, A–1
 - cylinder
 - definition, *Vol 2*, A–2
 - deallocating drives, *Vol 1*, 9–16
 - default format, *Vol 1*, 10–35
 - definition, *Vol 1*, 9–2
 - dismounting, *Vol 1*, 11–17
 - extents, *Vol 2*, A–1
 - definition, *Vol 1*, 9–3
 - file identification, *Vol 2*, A–4
 - files
 - See Disk files
 - Files–11
 - directory hierarchy, *Vol 2*, A–4
 - fragmentation of, *Vol 1*, 11–51; *Vol 2*, 16–34
 - I/O performance, *Vol 2*, 17–7
 - initializing, *Vol 1*, 11–15
 - mounting, *Vol 1*, 9–35, 11–16
 - in host-based shadow set, *Vol 1*, 11–45
 - organization
 - logical, *Vol 2*, A–1
 - physical, *Vol 2*, A–2
 - protecting, *Vol 1*, 9–21, 9–23
 - space
 - See Disk space
 - structure
 - See Disk structure
 - system
 - See System disks
 - terminology, *Vol 1*, 9–2
 - time limit, *Vol 1*, 9–75
 - track
 - definition, *Vol 2*, A–2
 - usage, *Vol 1*, 9–64
 - creating file, *Vol 1*, 9–73
 - UICs kept in quota file, *Vol 2*, A–9
 - user
 - See User disks
 - volumes

Disks

volumes (cont'd)

See Disk volumes

Disk space

See also Disk quotas

allocation by cluster, *Vol 2, A-1*

determining amount of free space, *Vol 2, 17-9*

managing, *Vol 1, 9-62 to 9-70*

purging files, *Vol 1, 9-67*

saving, *Vol 1, 9-67*

by moving page and swap files off the system disk, *Vol 2, 16-6*

by purging the operator log file, *Vol 1, 5-14*

by removing optional system files, *Vol 1, 5-1*

by storing minimal system dump information, *Vol 2, 16-12*

by using a selective system dump, *Vol 2, 16-4*

tracking use of, *Vol 2, 21-6*

Disk storage server, *Vol 2, 25-11*

Disk structure

analyzing errors, *Vol 1, 9-70*

Files-11, *Vol 2, A-3*

reporting and repairing errors, *Vol 1, 9-71*

Disk volumes

accessing files on, *Vol 1, 10-28*

access to public, *Vol 1, 9-14*

adding to an existing set, *Vol 1, 9-47*

adding volumes to volume sets, *Vol 1, 9-47*

analyzing disk structure errors, *Vol 1, 9-70*

analyzing media errors, *Vol 1, 9-80*

assigning logical name, *Vol 1, 5-11*

assigning volume label, *Vol 1, 5-11*

binding into volume sets, *Vol 1, 9-43*

characteristics

modifying, *Vol 1, 9-41*

console, *Vol 1, 10-40*

copying files from, *Vol 1, 10-35*

copying files to and from foreign volumes, *Vol 1, 10-40*

copying files to tape volumes, *Vol 1, 10-37*

creating Files-11 structure, *Vol 1, 9-16*

creating shadowed volume sets, *Vol 1, 9-46*

creating volume sets from, *Vol 1, 9-45*

definition, *Vol 2, A-1*

definitions, *Vol 1, 9-3*

disk quota operations, *Vol 1, 9-64*

dismounting, *Vol 1, 9-57*

displaying type, *Vol 1, 10-3*

file expiration dates

setting, *Vol 1, 9-68*

file-structured, *Vol 1, 9-28*

foreign, *Vol 1, 9-28*

handling error conditions, *Vol 1, 9-70*

initializing, *Vol 1, 9-16; Vol 2, A-1*

guidelines, *Vol 1, 9-19*

load balancing, *Vol 1, 9-15*

Disk volumes (cont'd)

modifying characteristics, *Vol 1, 9-41*

mounting, *Vol 1, 9-28, 9-34*

for page and swap files, *Vol 2, 16-26*

for queue database files, *Vol 1, 13-6*

in system startup, *Vol 1, 5-11*

early, *Vol 1, 5-12*

for page and swap files, *Vol 1, 5-6*

InfoServer, *Vol 2, 25-16*

MOUNT/ASSIST command, *Vol 1, 5-12*

special consideration about operator assistance, *Vol 1, 5-12*

mount verification, *Vol 1, 9-73*

performance, *Vol 1, 9-35*

physical loading, *Vol 1, 9-28*

private, *Vol 1, 9-15*

protecting, *Vol 1, 9-21*

public

See Public volumes

reading files from, *Vol 1, 10-28*

rebuilding, *Vol 1, 9-64*

removing before dismounting, *Vol 1, 9-58*

space

conserving, *Vol 1, 9-63*

verification, *Vol 1, 9-72*

write-locked

dismounting, *Vol 1, 9-58*

write-locking, *Vol 1, 9-74*

writing files to, *Vol 1, 10-35*

DISMOUNT command

See also Dismounting

canceling mount verification, *Vol 1, 9-77*

for a single tape volume, *Vol 1, 9-59*

for foreign volumes, *Vol 1, 9-59*

for single volume in volume set, *Vol 1, 9-59*

for single volume or volume set, *Vol 1, 9-57*

overriding automatic unloading of volume, *Vol 1, 9-59*

tape, *Vol 1, 10-38*

Dismounting

See also DISMOUNT command

a backup volume, *Vol 1, 11-17*

a disk with DECdtm transaction logs, *Vol 2, 27-17*

system disk, *Vol 1, 9-57*

volumes

conditions preventing, *Vol 1, 9-57*

foreign, *Vol 1, 9-59*

in an OpenVMS Cluster system, *Vol 1, 9-60*

with cached information, *Vol 1, 9-57*

with open files, *Vol 1, 9-57*

volume sets, *Vol 1, 9-59*

Displaying

characteristics assigned to a queue, *Vol 1, 14-31*

Displaying (cont'd)

- defined characteristics, *Vol 1*, 14–31
- defined forms, *Vol 1*, 14–48
- forms assigned to queues, *Vol 1*, 14–48
- information about queue managers, *Vol 1*, 13–6
- information about queues, *Vol 1*, 14–58
- system parameters

See Showing system parameters

Distributed network management

- concepts, *Vol 2*, 23–5, 23–10

Distributed Queuing System

See DQS

Distributing system work load, *Vol 2*, 17–3

Distribution kit

- startup files included on, *Vol 1*, 5–2

DNA (Digital Network Architecture), *Vol 2*, 23–9

Phase V

See DECnet-Plus

DO command

- for managing an OpenVMS Cluster system, *Vol 2*, 22–17

in SYSMAN, *Vol 1*, 2–18

DOS–11 tape volumes

- file transfers with, *Vol 1*, 10–40
- format conversions for, *Vol 1*, 10–40

DOSD (dump off system disk), *Vol 2*, 16–14, 16–15, 16–20

- requirements, *Vol 2*, 16–18

Downline loading, *Vol 2*, 25–2

DQS (Distributed Queuing Service)

- distributed printing, *Vol 1*, 14–13

Drive

- definition, *Vol 1*, 9–2

Drivers

See Device drivers

DSA device naming, *Vol 1*, 8–1

DSI (DIGITAL System Identifier)

- ISO 9660 media protection, *Vol 1*, 9–25
- mount option, *Vol 1*, 9–24

DSI keyword

- with MOUNT/PROTECTION command, *Vol 1*, 9–30

DTE (data terminal equipment), *Vol 1*, 8–27, 8–28

Dual-architecture OpenVMS Cluster systems

- installing images, *Vol 2*, 22–18
- example, *Vol 2*, 22–18

Dual format

- ISO 9660 CD and ODS, *Vol 1*, 9–4

DUMPBUG system parameter, *Vol 2*, 16–4

Dump file information

- saving automatically, *Vol 2*, 16–19

Dump files

- See also Crash dumps; Error log dump file; System dump file; System failures
- creating
 - with AUTOGEN, *Vol 2*, 16–31
- definition, *Vol 2*, 16–2

Dump files (cont'd)

- information captured in, *Vol 2*, 16–2
- location

specifying, *Vol 2*, 16–31

- overwriting of, *Vol 2*, 16–2

sizes

changing, *Vol 2*, 16–28, 16–32

specifying, *Vol 2*, 16–31

system

See System dump files

- tasks for managing, *Vol 2*, 16–1

DUMPFILe symbol, *Vol 2*, 16–31

DUMPFILe_DEVICE symbol, *Vol 2*, 16–29

Dump off system disk (DOSD), *Vol 2*, 16–15

DUMPSTYLE system parameter, *Vol 2*, 16–4, 16–7, 16–12

Dynamic load balancing, *Vol 2*, 28–2

Dynamic system parameters, *Vol 2*, 15–2, 15–3

See also System parameters

E

EDIT keypad function, *Vol 2*, 22–7

Emergency system shutdown

- with console commands, *Vol 1*, 4–22

with OPCCRASH, *Vol 1*, 4–22, 4–31

Emergency system startup

- with default system parameters, *Vol 1*, 4–9

without startup and login procedures, *Vol 1*, 4–10

without the UAF, *Vol 1*, 4–11

/ENABLE qualifier

SET RIGHTS_LIST command, *Vol 1*, 10–14

Enabling autostart, *Vol 1*, 14–4, 14–17

in startup command procedure, *Vol 1*, 14–17

recommended use, *Vol 1*, 14–18

using ENABLE AUTOSTART/QUEUES command, *Vol 1*, 14–17

END phase

of system startup, *Vol 1*, 5–19

ENQLM process limit, *Vol 1*, 7–45

EOT markers, *Vol 1*, 9–11

continuing to copy after, *Vol 1*, 10–39

Erasing blocks, *Vol 1*, 9–73

ERLBUFFERPAGES system parameter, *Vol 2*, 16–7

ERRFMT

See Error Formatter

ERRFMT process, *Vol 2*, 20–2

See also Error log files; Error Log utility; Error logging

creation during system startup, *Vol 1*, 5–5

restarting, *Vol 2*, 20–3

sending mail when it is deleted, *Vol 2*, 20–4

writes to ERRLOG.SYS file, *Vol 2*, 20–2

ERRLOGDUMP symbol, *Vol 2*, 16–31

Error checking
 in SYSTARTUP_VMS.COM, *Vol 1*, 5–10
 in system parameter files, *Vol 2*, 15–19

Error Formatter
 changing mail user name, *Vol 2*, 20–5
 description, *Vol 2*, 20–2
 disabling mail, *Vol 2*, 20–5
 enabling mail, *Vol 2*, 20–5
 notifying user with mail message, *Vol 2*, 20–4
 using to send mail, *Vol 2*, 20–4

ERRORLOG.DMP file
 See Error log dump file

ERRORLOGBUFFERS system parameter, *Vol 2*, 16–7

Error log dump files
 created by AUTOGEN, *Vol 2*, 16–3
 creating with AUTOGEN, *Vol 2*, 16–31
 definition, *Vol 2*, 16–2
 installing
 automatically, *Vol 2*, 16–3
 rebooting after creation, *Vol 2*, 16–34
 sizes, *Vol 2*, 16–31
 calculating manually, *Vol 2*, 16–9
 calculating with AUTOGEN, *Vol 2*, 16–3
 manipulating with AUTOGEN, *Vol 2*, 16–30

Error log files
 created by ERRFMT process, *Vol 2*, 20–2
 events reported in, *Vol 2*, 20–5
 logical name defining location, *Vol 1*, 5–8
 maintaining, *Vol 2*, 20–4
 moving to reduce system disk I/O, *Vol 2*, 17–7
 producing full report, *Vol 2*, 20–7
 reporting on using DECEvent, *Vol 2*, 20–9
 SYSPRV privilege to access, *Vol 2*, 20–7

Error logging
 See also ERRFMT process; Error log files; Error Log utility
 description, *Vol 2*, 20–2
 facility, *Vol 2*, 20–2
 reports produced, *Vol 2*, 20–3

Error Log Report Formatter (ERF), *Vol 2*, 20–2

Error log reports
 See Error Log utility, reports

Error Log utility (ERROR LOG)
 See also ERRFMT process; Error logging; Error log files
 ANALYZE/ERROR_LOG command, *Vol 2*, 20–7
 definition, *Vol 2*, 20–5
 error logging, *Vol 2*, 20–2
 relationship to UETP, *Vol 2*, 19–3, 19–17, 19–29
 reporting on error log file, *Vol 2*, 20–5
 reports
 excluding unknown entries, *Vol 2*, 20–8
 formats, *Vol 2*, 20–8

Error Log utility (ERROR LOG)
 reports (cont'd)
 printing, *Vol 2*, 20–7
 privileges required, *Vol 2*, 20–7
 specifying display device, *Vol 2*, 20–8
 specifying events and times, *Vol 2*, 20–8
 types, *Vol 2*, 20–6
 uses, *Vol 2*, 20–6
 translating error log files, *Vol 2*, 20–5

Error messages, *Vol 1*, 2–5
 See also Help Message utility; Messages

Error options
 for fatal BACKUP errors, *Vol 1*, 11–68

Errors
 analyzing disk structure, *Vol 1*, 9–70
 analyzing error reports, *Vol 2*, 20–2
 analyzing media, *Vol 1*, 9–80
 disk read
 if returned when booting, *Vol 1*, 4–18
 disk structure
 repairing, *Vol 1*, 9–71
 reporting, *Vol 1*, 9–71
 error log file, *Vol 2*, 20–2
 error logging facility, *Vol 2*, 20–2
 handling on disk volumes, *Vol 1*, 9–70
 machine check
 if returned when booting, *Vol 1*, 4–18
 mounting disk, *Vol 1*, 9–29

Errors during UETP, *Vol 2*, 19–21
 diagnosing, *Vol 2*, 19–16
 sources of, *Vol 2*, 19–17

ESS\$LASTDRIVER device driver, *Vol 2*, 25–9, 25–15
 controlling and diagnosing, *Vol 2*, 25–9, 25–11

ESS\$STARTUP.COM command procedure, *Vol 2*, 25–11, 25–15
 invoking in system startup, *Vol 2*, 25–11

Ethernet
 adapters
 specifying number of in MODPARAMS.DAT file, *Vol 2*, 15–21
 controllers
 problems with multiple controllers, *Vol 2*, 25–12
 defining a remote node for UETP, *Vol 2*, 19–20
 device drivers, *Vol 2*, 24–2
 preparing for UETP, *Vol 2*, 19–8

Event classes
 displaying those being audited, *Vol 2*, 20–28

Event handling
 device driver used in, *Vol 1*, 8–9

EXCHANGE/NETWORK command
 using to transfer files, *Vol 1*, 10–34, 10–40

Exchange utility (EXCHANGE)
 using to copy files, *Vol 1*, 10–34
 using to transfer information, *Vol 1*, 10–40

- Executable images, *Vol 2*, 17–17, 17–22, 17–23
- Execute access
 - for disk directory files, *Vol 1*, 10–26
 - for disk files, *Vol 1*, 10–21
 - granting through protection codes, *Vol 1*, 12–8
- Execute procedure (@) command, *Vol 2*, 22–16
- Executing job status, *Vol 1*, 14–66, 14–69
- Execution modes
 - startup procedures, *Vol 1*, 5–20
 - BATCH, *Vol 1*, 5–20
 - changing, *Vol 1*, 5–21
 - DIRECT, *Vol 1*, 5–20
 - SPAWN, *Vol 1*, 5–20
 - specifying, *Vol 1*, 5–21
- Execution queues
 - activating autostart, *Vol 1*, 14–4, 14–16
 - creating, *Vol 1*, 14–15
 - relationship to generic queues, *Vol 1*, 14–3
 - starting
 - autostart, *Vol 1*, 14–4, 14–17
 - in system startup, *Vol 1*, 5–12
 - nonautostart, *Vol 1*, 14–18
- Executive mode
 - calling images running in, *Vol 2*, 17–19, 17–22
 - logical names, *Vol 2*, 17–23
 - recommended use for logical names, *Vol 1*, 5–9
- Expanding system libraries, *Vol 2*, 17–8
- Expiration date, *Vol 1*, 7–45
 - field, *Vol 1*, 10–29
 - checking, *Vol 1*, 10–33
 - file, *Vol 1*, 9–68
 - tape file system checks, *Vol 1*, 10–33
- Expiration time, *Vol 1*, 7–45
- Extended attribute records
 - See XARs
- Extended File Cache, *Vol 2*, 18–1
- Extended File Specifications
 - backing up ODS-5 volumes, *Vol 1*, 10–9
 - BACKUP limitations, *Vol 1*, 9–10
 - BACKUP support of, *Vol 1*, 11–6
 - case blindness, *Vol 1*, 10–4
 - case preservation, *Vol 1*, 10–4
 - CONDENSED file names, *Vol 1*, 10–5
 - converting existing volumes, *Vol 1*, 9–36
 - converting from ODS-1 to ODS-2, *Vol 1*, 9–39
 - converting from ODS-5 to ODS-2, *Vol 1*, 9–39
 - deep directory levels, *Vol 1*, 10–2
 - DID (Directory ID)
 - displaying, *Vol 1*, 10–5
 - displaying volume structure, *Vol 1*, 10–3
 - enabling, *Vol 1*, 9–10
 - enabling parsing, *Vol 1*, 10–2
 - extended file names, *Vol 1*, 9–8
 - EXTENDED parse style, *Vol 1*, 10–6
 - file name parsing
 - related error messages, *Vol 1*, 10–6
 - file names
 - avoiding incompatibility, *Vol 1*, 10–6

- Extended File Specifications
 - file names (cont'd)
 - implicit output, *Vol 1*, 10–7
 - file naming, *Vol 1*, 10–1
 - initializing a new volume, *Vol 1*, 9–20
 - ISO Latin-1 character set, *Vol 1*, 10–2
 - mixed-architecture support, *Vol 1*, 9–9
 - mixed-version cluster support, *Vol 1*, 9–9
 - mixed-version support, *Vol 1*, 9–9
 - network support, *Vol 1*, 9–10
 - ODS-5 extended file names
 - with batch command file, *Vol 1*, 10–8
 - ODS-5 support, *Vol 1*, 10–2
 - OpenVMS Version 7.2 support, *Vol 1*, 9–9
 - submitting a batch command file, *Vol 1*, 10–8
 - system support, *Vol 1*, 9–9
 - traditional file names, *Vol 1*, 9–8
 - TRADITIONAL parse style, *Vol 1*, 10–2
 - Unicode (UCS-2) character set, *Vol 1*, 10–2
 - usage notes, *Vol 1*, 10–3
 - user disk support, *Vol 1*, 9–9
 - using DCL with, *Vol 1*, 10–2
 - using RMS features, *Vol 1*, 10–2
 - VAX system limitations, *Vol 1*, 9–9
 - volume structures
 - converting existing volumes to ODS-5, *Vol 1*, 9–36
 - migrating to another format, *Vol 1*, 9–39
 - Windows 95 style or Windows NT style file names, *Vol 1*, 9–8
- Extensions
 - file
 - See File extensions
- Extents
 - definition, *Vol 2*, A–1
 - disk
 - definition, *Vol 1*, 9–3
 - index file
 - definition, *Vol 2*, A–7
- Extracting software product release notes, *Vol 1*, 3–25

F

- F\$GETJPI lexical function
 - getting information about vector processing, *Vol 2*, 28–9
- F\$GETQUI lexical function, *Vol 1*, 14–59
- F\$GETSYI lexical function
 - getting information about vector processing, *Vol 2*, 28–9
- Failover list
 - for an autostart queue
 - specifying, *Vol 1*, 14–16
 - for queue manager, *Vol 1*, 13–3, 13–10
 - insufficient, *Vol 1*, 13–19
 - specifying, *Vol 1*, 13–9

- Failovers
 - See also Failover list
 - of queue manager, *Vol 1*, 13–3, 13–19
 - forcing, *Vol 1*, 13–10
 - of queues, *Vol 1*, 14–4, 14–16
- FDDI (Fiber Distributed Data Interface)
 - device drivers, *Vol 2*, 24–2
- FDDI controllers
 - problems with multiple controllers, *Vol 2*, 25–12
- Feedback
 - See AUTOGEN feedback
- FID (file identification)
 - See File identification
- FIELD account
 - initial modification, *Vol 1*, 7–9
 - in UAF, *Vol 1*, 7–7
- Field of data
 - in SHOW CLUSTER, *Vol 2*, 22–5
 - removing, *Vol 2*, 22–11
- Field Service accounts
 - in UAF, *Vol 1*, 7–7
- File access
 - disk, *Vol 1*, 10–28
 - levels allowed in restore operations, *Vol 1*, 11–31
 - listing number of concurrent, *Vol 2*, 17–7
 - tape, *Vol 1*, 10–28, 10–29
- File attributes, *Vol 1*, 10–18
- File banner pages, *Vol 1*, 14–46
 - See also Job banner pages
- File-based autoconfiguration, *Vol 1*, 8–11
- File extensions
 - effect on system performance, *Vol 2*, 17–7
 - specifying size, *Vol 1*, 9–30; *Vol 2*, 17–7
 - system parameter controlling, *Vol 2*, 17–7
- File formats
 - use with BACKUP, *Vol 1*, 11–10
- File fragmentation
 - of page and swap files, *Vol 2*, 16–34
- File headers
 - index file, *Vol 2*, A–7
 - contents, *Vol 2*, A–7
 - extension, *Vol 2*, A–8
 - primary, *Vol 2*, A–8
- File identification
 - file number, *Vol 2*, A–4
 - Files–11, *Vol 2*, A–4
 - file sequence number (SEQ), *Vol 2*, A–4
 - relative volume number (RVN), *Vol 2*, A–4
- File Log
 - VMSINSTAL.COM option, *Vol 1*, 3–13
- File names
 - BACKUP processes extended character sets, *Vol 1*, 11–6
 - displaying with EXPANDED style, *Vol 1*, 10–5
- File names (cont'd)
 - displaying with Extended File Specifications
 - CONDENSED style, *Vol 1*, 10–5
 - Extended File Specifications, *Vol 1*, 10–2
 - Extended File Specifications styles
 - extended, *Vol 1*, 9–8
 - traditional, *Vol 1*, 9–8
- File protection, *Vol 1*, 10–19
 - SYSDUMP.DMP file, *Vol 2*, 16–5
- Files
 - See also Files-11 On-Disk Structure; Parameter files
 - accessing
 - See File access
 - adding and removing optional, *Vol 1*, 5–1
 - attributes
 - accessing, *Vol 1*, 10–32
 - backing up, *Vol 1*, 11–24
 - caching, *Vol 2*, 18–1
 - comparing with BACKUP, *Vol 1*, 11–27
 - copying
 - from BACKUP, *Vol 1*, 11–24
 - from disk to standard-labeled volumes, *Vol 1*, 10–35
 - from disk volumes, *Vol 1*, 10–35
 - to tape, *Vol 1*, 10–37
 - to tape volumes, *Vol 1*, 10–37
 - creating, *Vol 1*, 9–5
 - detected bad block (DBBF), *Vol 1*, 9–80
 - expiration dates on, *Vol 1*, 9–68
 - for AUTOGEN feedback, *Vol 2*, 15–11
 - limiting number of versions, *Vol 1*, 9–68
 - logging activity during installation, *Vol 1*, 3–13
 - lost
 - recovering, *Vol 1*, 9–72
 - modifying characteristics, *Vol 1*, 10–24
 - naming
 - on Files–11 volume, *Vol 2*, A–7
 - nonstandard format
 - DCL commands with, *Vol 1*, 10–16, 10–28
 - on public volumes, *Vol 1*, 9–14
 - open during backup, *Vol 1*, 11–32, 11–66
 - overwriting, *Vol 1*, 10–29
 - PCF, *Vol 1*, 3–19
 - private volumes, *Vol 1*, 9–15
 - privileges, *Vol 1*, 10–19
 - public, *Vol 1*, 9–14
 - purging to save disk space, *Vol 1*, 9–68
 - recovering lost, *Vol 1*, 9–72
 - reserved, *Vol 2*, A–4
 - list of, *Vol 1*, 9–5
 - restoring
 - from BACKUP, *Vol 1*, 11–29, 11–30
 - from image backup, *Vol 1*, 11–45
 - from incremental backup, *Vol 1*, 11–47
 - retrieving information from, *Vol 1*, 10–16
 - security
 - using protection codes, *Vol 1*, 10–21

Files (cont'd)

- system
 - moving to reduce system disk I/O, *Vol 2*, 17–8
- tape
 - See Tape files
- tape volumes, *Vol 1*, 10–20
 - writing to files on, *Vol 1*, 10–33
- transferring across network, *Vol 1*, 10–40
- versions
 - limiting number of, *Vol 1*, 9–68, 10–27
 - VMSMAIL_PROFILE.DATA, *Vol 1*, 7–41
 - writing to a CD-R or CD-RW disk, *Vol 1*, 10–43
- Files-11 data caches, *Vol 2*, 18–1
- Files-11 metadata caches, *Vol 2*, 18–2
- Files-11 On-Disk Structure
 - block
 - definition, *Vol 2*, A–1
 - comparison of ODS Levels 1, 2, and 5, *Vol 2*, A–10
 - creating structure, *Vol 1*, 9–16
 - disk save set, *Vol 1*, 11–8
 - file identification, *Vol 2*, A–4
 - Level 1, *Vol 1*, 10–35
 - directory hierarchy and structure, *Vol 2*, A–4
 - Level 2, *Vol 1*, 10–35
 - assigning disk quotas, *Vol 1*, 9–65
 - directory hierarchy and structure, *Vol 2*, A–4
 - Level 5
 - assigning disk quotas, *Vol 1*, 9–65
 - directory hierarchy and structure, *Vol 2*, A–4
 - master file directory (MFD), *Vol 2*, A–4
 - sector, *Vol 2*, A–2
 - structure, *Vol 2*, A–3
 - Level 1, *Vol 2*, A–4
 - Level 2, *Vol 2*, A–4
 - Level 5, *Vol 2*, A–4
 - terminology, *Vol 2*, A–2
 - UIC, *Vol 2*, A–4
 - using with Exchange utility (EXCHANGE)
 - to transfer data, *Vol 1*, 10–40
- File specifications
 - ANSI, *Vol 1*, 10–31
 - for installing images, *Vol 2*, 17–24
 - OpenVMS extended, *Vol 1*, 10–31
 - standard, *Vol 1*, 10–31
- File structures
 - disk and CD-ROM, *Vol 1*, 9–4
 - options on OpenVMS systems, *Vol 1*, 9–4
- File system
 - controlling data caching, *Vol 2*, 18–1
- File versions
 - relative input
 - BACKUP supports, *Vol 1*, 11–6

File windows

- mapping pointers for, *Vol 1*, 9–31

FILLM process limit, *Vol 1*, 7–45

- value for efficient backups, *Vol 1*, 11–11

Fix-ups

- performing at installation time, *Vol 2*, 17–20

Flag pages, *Vol 1*, 14–39

- file, *Vol 1*, 14–43
- job, *Vol 1*, 14–42

Foreign volumes

- See Volumes, foreign

Formats

- CD-ROM on disc, *Vol 1*, 9–4
- converting software kits, *Vol 1*, 3–30
- High Sierra, *Vol 1*, 9–3

Formatting

- volumes, *Vol 1*, 9–18

Form feeds

- controlling page overflow with, *Vol 1*, 14–52
- inserting automatically in output jobs, *Vol 1*, 14–21
- suppressing, *Vol 1*, 14–52

Forms

- assigning default for a queue, *Vol 1*, 14–50
- associating with jobs and queues, *Vol 1*, 14–47
- changing DEFAULT, *Vol 1*, 14–48
- commands used with, *Vol 1*, 14–47
- controlling line overflow with, *Vol 1*, 14–51
- controlling page width, length, and margin size with, *Vol 1*, 14–47
- controlling paper stock with, *Vol 1*, 14–47
- creating, *Vol 1*, 14–49
- default, *Vol 1*, 14–50
- DEFAULT, *Vol 1*, 14–48
- deleting, *Vol 1*, 14–51
 - problems, *Vol 1*, 14–76
- description, *Vol 1*, 14–46
- displaying
 - defined forms, *Vol 1*, 14–48
 - forms assigned to queues, *Vol 1*, 14–48
- formatting jobs with, *Vol 1*, 14–47
- mounting, *Vol 1*, 14–51
- preprinted
 - aligning, *Vol 1*, 14–70, 14–71
 - procedure for using, *Vol 1*, 14–47
 - specifying setup modules with, *Vol 1*, 14–47
 - specifying sheet-fed paper with, *Vol 1*, 14–47

- Fragmentation of disks, *Vol 1*, 11–51
- Full backup
- See Image backup
- Full names
- DECnet-Plus
 - syntax, *Vol 2*, 23–10
- definition, *Vol 2*, 23–9

G

- GBLPAGES system parameter, *Vol 2*, 17–21
- GBLSECTIONS system parameter, *Vol 2*, 17–21
- Generic queues
 - batch, *Vol 1*, 14–4
 - recommended use, *Vol 1*, 14–6
 - creating, *Vol 1*, 14–19
 - description, *Vol 1*, 14–3, 14–4
 - in an OpenVMS Cluster environment, *Vol 1*, 14–3
 - output, *Vol 1*, 14–4, 14–10, 14–11
 - recommended use, *Vol 1*, 14–10
 - relationship to execution queues, *Vol 1*, 14–3
- Get Save Set
 - VMSINSTAL.COM option, *Vol 1*, 3–11, 3–12
- Getting help, *Vol 1*, 3–26
- Global pages, *Vol 2*, 17–21
- Global sections, *Vol 2*, 17–21
- Group numbers
 - modifying, *Vol 2*, 22–16
- Groups
 - accounting, *Vol 2*, 21–5
- Group users (security category), *Vol 1*, 12–8
- Group volumes
 - definition, *Vol 1*, 9–14
- GRPPRV privilege
 - giving rights of system user, *Vol 1*, 12–8

H

- Halting system
 - waiting until after system dump file is written, *Vol 2*, 16–2
- Hard links, *Vol 1*, 10–47
 - enabling, *Vol 1*, 10–48
- Hardware
 - booting problem, *Vol 1*, 4–18
 - importance of sufficient capacity for system performance, *Vol 2*, 17–5
- Header labels
 - on tape files, *Vol 1*, 9–12, 10–30
 - reading attributes of, *Vol 1*, 10–32
- Header resident images, *Vol 2*, 17–18, 17–20
- Help
 - POLYCENTER Software Installation utility, *Vol 1*, 3–26
- HELP command
 - affected during layered product installation, *Vol 1*, 3–28
- Help Message utility (MSGHLP), *Vol 1*, 5–25
 - accessing \$STATUS values for uninstalled messages, *Vol 1*, 5–26
 - adding .MSGHLP\$DATA files to the database, *Vol 1*, 5–28
 - adding comments to the database, *Vol 1*, 5–30
 - adding messages to the database, *Vol 1*, 5–32

- Help Message utility (MSGHLP) (cont'd)
 - changing Compaq supplied data, *Vol 1*, 5–31
 - compressing the database after deletions, *Vol 1*, 5–30
 - creating databases for different user groups, *Vol 1*, 5–28
 - customizing the database, *Vol 1*, 5–25
 - default database, *Vol 1*, 5–28
 - deleting Compaq supplied messages from the database, *Vol 1*, 5–29
 - message section files, *Vol 1*, 5–26
 - order searched, *Vol 1*, 5–27
 - search path of database files, *Vol 1*, 5–28
 - /SECTION_FILE qualifier usage, *Vol 1*, 5–26
- Hierarchical storage controller devices
 - See HSC devices
- High Sierra format
 - definition, *Vol 1*, 9–4
 - on CD-ROM, *Vol 1*, 9–3
- High-water marking
 - disabling for system performance, *Vol 2*, 17–6
- Holding a job, *Vol 1*, 14–67
- Holding job status, *Vol 1*, 14–66, 14–73
- Home blocks, *Vol 1*, 9–73
 - in index file, *Vol 2*, A–6
- Host-based shadow set
 - mounting disks, *Vol 1*, 11–45
- HSC devices
 - configuring during system startup, *Vol 1*, 8–7
 - disabling configuration during system startup, *Vol 1*, 8–10
- Hyphen (-), restriction on use of, *Vol 1*, 3–14

I

- I/O (input/output)
 - reducing on system disk, *Vol 2*, 17–7
- ICC (Intra-Cluster communications), *Vol 1*, 12–9
- ICC security objects, *Vol 1*, 12–9
- Identification record
 - ANALYZE/DISK_STRUCTURE, *Vol 1*, 9–70
- Identifier fields
 - file, *Vol 1*, 10–30
 - volume, *Vol 1*, 9–53
- Identifiers
 - Environmental, *Vol 1*, 12–11
 - General, *Vol 1*, 12–11
 - system-defined, *Vol 1*, 12–11
 - types, *Vol 1*, 12–11
 - UIC, *Vol 1*, 12–11
- Idle queue status, *Vol 1*, 14–59
- Image backup
 - command format for disks, *Vol 1*, 11–34
 - command format for tapes, *Vol 1*, 11–33
 - definition, *Vol 1*, 11–3
 - restoring files from, *Vol 1*, 11–45
 - to disk, *Vol 1*, 11–34
 - to tape, *Vol 1*, 11–33, 11–34

Image Registry facility, *Vol 1*, 5–23

Images

See also Known images

application, *Vol 2*, 17–21

concurrent access by multiple users, *Vol 2*, 17–19

definition, *Vol 2*, 17–17

determining frequency of use, *Vol 2*, 17–19, 17–25

executable, *Vol 2*, 17–17, 17–22, 17–23

execute-only, *Vol 2*, 17–23

header resident, *Vol 2*, 17–18, 17–20

installing

See also Known images

application, *Vol 2*, 17–21

effect on RUN command, *Vol 2*, 17–18

in system startup, *Vol 1*, 5–4, 5–13; *Vol 2*, 17–18

reasons for, *Vol 2*, 17–17

to improve system performance, *Vol 2*, 17–7

with shared address data, *Vol 2*, 17–20, 17–21

known, *Vol 2*, 17–18

See also Known images

linkable, *Vol 2*, 17–17

permanently open, *Vol 2*, 17–18, 17–20

privileged, *Vol 2*, 17–19, 17–22

security caution, *Vol 2*, 17–22

shareable, *Vol 2*, 17–22

protected, *Vol 2*, 17–19, 17–22

protecting installed, *Vol 2*, 17–23

resident (Alpha), *Vol 2*, 17–19

running in protected modes, *Vol 2*, 17–19, 17–22

shareable, *Vol 2*, 17–19, 17–21, 17–22

assigning logical names for, *Vol 2*, 17–25

system-provided, *Vol 2*, 17–21

system version dependent

registering, *Vol 1*, 5–23

user-level

calling of protected code, *Vol 2*, 17–19, 17–22

version checking, *Vol 1*, 5–23

writable, *Vol 2*, 17–19

Incremental backup

command format for disks, *Vol 1*, 11–37

command format for tapes, *Vol 1*, 11–36

definition, *Vol 1*, 11–3

restoring files from, *Vol 1*, 11–47

to disk, *Vol 1*, 11–36

to tape, *Vol 1*, 11–35

INDEXF.SYS file

See Index files

Index files, *Vol 2*, A–4 to A–8

alternate file header, *Vol 2*, A–6

backup home block, *Vol 2*, A–6

Index files (cont'd)

backup index file header, *Vol 2*, A–6

bitmap, *Vol 2*, A–6

boot block, *Vol 2*, A–5, A–6

bootstrap image, *Vol 2*, A–6

definition, *Vol 2*, A–5

file headers, *Vol 2*, A–6, A–8

file number, *Vol 2*, A–4

home block, *Vol 2*, A–5, A–6

INDEXF.SYS, *Vol 2*, A–5

in volume sets, *Vol 1*, 9–45

reserved file, *Vol 2*, A–5

InfoServer

See also InfoServer Client for OpenVMS

automatic service, *Vol 2*, 25–4

availability, *Vol 2*, 25–4

backing up system disks, *Vol 1*, 11–64

Client

and DECnet, *Vol 2*, 25–11

commands, *Vol 2*, 25–7

console terminal, *Vol 2*, 25–6

determining default server name, *Vol 2*, 25–7

downline loading with, *Vol 2*, 25–2

failover, *Vol 2*, 25–4

functions, *Vol 2*, 25–1

Help facility, *Vol 2*, 25–8

load balancing, *Vol 2*, 25–6

local connections, *Vol 2*, 25–6

management session, ending, *Vol 2*, 25–7

mounting devices

in system startup, *Vol 2*, 25–16

multicast address feature, *Vol 2*, 25–6

password, *Vol 2*, 25–7

protocols, *Vol 2*, 25–5

quick access to duplicate services for clients, *Vol 2*, 25–4

relationship to client systems, *Vol 2*, 25–2

remote connections, *Vol 2*, 25–6

removing media, *Vol 2*, 25–4

server name, *Vol 2*, 25–7

service disconnection, *Vol 2*, 25–4

setting up in system startup, *Vol 2*, 25–11

software

client failure to start, *Vol 2*, 25–12

client transport (LAST) problems, *Vol 2*, 25–12

starting client for OpenVMS, *Vol 2*, 25–11

updating from the Software Products

Library, *Vol 2*, 25–3

starting a session, *Vol 2*, 25–6

starting Client, *Vol 2*, 25–11

support for X terminals, *Vol 2*, 25–5

system overview, *Vol 2*, 25–1

virtual device server, *Vol 2*, 25–1

virtual device units, *Vol 2*, 25–15

X terminal client, *Vol 2*, 25–5

- InfoServer Client for OpenVMS
 - components, *Vol 2*, 25–9
 - functions, *Vol 2*, 25–9
 - mounting devices
 - in system startup, *Vol 1*, 5–13
 - setting up in system startup, *Vol 1*, 5–13
 - software, *Vol 2*, 25–9
 - starting automatically, *Vol 2*, 25–11
 - startup restrictions, *Vol 2*, 25–14
- Initialization files
 - creating, *Vol 2*, 22–15
 - establishing SHOW CLUSTER reports, *Vol 2*, 22–6
 - SHOW CLUSTER
 - creating, *Vol 2*, 22–14
 - SHOW_CLUSTER\$INIT, *Vol 2*, 22–14, 22–15
 - use with SYSMAN, *Vol 1*, 2–19
- Initialization of system
 - in a multiprocessing system, *Vol 2*, 28–2
- INITIALIZE/QUEUE command, *Vol 1*, 14–18
 - activating autostart queues, *Vol 1*, 14–16
 - assigning a default form, *Vol 1*, 14–50
 - assigning characteristics, *Vol 1*, 14–31
 - canceling characteristics, *Vol 1*, 14–31
 - controlling page overflow, *Vol 1*, 14–52
 - creating a generic queue, *Vol 1*, 14–19
 - mounting a form, *Vol 1*, 14–51
 - setting UIC-based protection on queues, *Vol 1*, 14–24
 - specifying autostart information, *Vol 1*, 14–16
 - specifying banner pages, *Vol 1*, 14–46
 - specifying job processing options, *Vol 1*, 14–35
 - specifying queue options, *Vol 1*, 14–20
 - specifying reset modules, *Vol 1*, 14–53
 - starting a nonautostart queue, *Vol 1*, 14–18
- INITIALIZE/STRUCTURE_LEVEL=5 command, *Vol 1*, 9–20
- INITIALIZE command
 - See also INITIALIZE/QUEUE command;
 - Initializing, volumes
 - creating volume identifiers for continuation volumes, *Vol 1*, 9–54
 - disk volumes, *Vol 1*, 9–17
 - for formatting page and swap file disks during system startup, *Vol 1*, 5–6
 - mounting volume sets, *Vol 1*, 9–54
 - qualifiers, *Vol 1*, 9–18
 - setting device protection, *Vol 1*, 9–23
 - for tape volumes, *Vol 1*, 10–28, 10–38
 - specifying logical names, *Vol 1*, 10–38
 - to format and write label to volume, *Vol 1*, 9–18
- Initializing
 - queues, *Vol 1*, 14–15
 - See also INITIALIZE/QUEUE command
 - volumes
 - See also INITIALIZE command
 - assisting users, *Vol 1*, 9–21
- Initializing
 - volumes (cont'd)
 - disk volumes, *Vol 1*, 9–16, 9–19
 - ODS-5 format, *Vol 1*, 9–20
 - results of, *Vol 1*, 11–14
 - tape volumes, *Vol 1*, 10–20
- INITIAL phase of system startup, *Vol 1*, 5–4, 5–18
- Initial System Load
 - See ISL
- Input an existing PCF, *Vol 1*, 3–26
- Input files
 - relative versions
 - BACKUP supports, *Vol 1*, 11–6
- Input symbiont for card reader
 - running interactively, *Vol 1*, 8–40
- Installation and upgrade information
 - Alpha only
 - problems and restrictions
 - tuning BAP system parameters, *Vol 2*, 15–37
- Installation procedures
 - See also Installing software; VMSINSTALL.COM command procedure
 - completing, *Vol 1*, 3–9
- INSTALL command
 - in SYSGEN, *Vol 2*, 16–25
 - in system startup, *Vol 1*, 5–7
- Installed files
 - See Known images
- Installing images, *Vol 2*, 17–24
 - See also Install utility; Known images
 - effect on RUN command, *Vol 2*, 17–18
 - in SYSTARTUP_VMS.COM, *Vol 1*, 5–13; *Vol 2*, 17–18
 - reasons for, *Vol 2*, 17–17
 - to improve system performance, *Vol 2*, 17–7, 17–18
- Installing page and swap files
 - in system startup, *Vol 1*, 5–6; *Vol 2*, 16–6, 16–26, 16–27
 - with AUTOGEN, *Vol 2*, 16–29
 - with SYPAGSWPFILES.COM command procedure, *Vol 2*, 16–26
 - with SYSGEN, *Vol 2*, 16–25
- Installing software
 - See also Installation procedure
 - as a batch job, *Vol 1*, 3–29
 - completing procedure, *Vol 1*, 3–9
 - creating a new PCF during, *Vol 1*, 3–26
 - deinstalling software, *Vol 1*, 3–31
 - logging file activity, *Vol 1*, 3–13
 - managing installed software, *Vol 1*, 3–31
 - multiple products at once, *Vol 1*, 3–25
 - noncompliant products, *Vol 1*, 3–23
 - on alternate disk, *Vol 1*, 3–9
 - preparing for installation, *Vol 1*, 3–3

- Installing software (cont'd)
 - reconfiguring product options, *Vol 1*, 3–29
 - removing installed software, *Vol 1*, 3–31
 - responding to questions during installation, *Vol 1*, 3–26
 - specifying source and destination locations, *Vol 1*, 3–25
 - supplying answers from a PCF, *Vol 1*, 3–26
 - tracking dependencies, *Vol 1*, 3–24
 - using a PCF for consistency, *Vol 1*, 3–19
 - Install utility (INSTALL)
 - See also Installing images; Known images and version checking, *Vol 1*, 5–23
 - determines content of address data section, *Vol 2*, 17–21
 - determining number frequency of image use, *Vol 2*, 17–19, 17–25
 - improving system performance, *Vol 2*, 17–7, 17–17, 17–18
 - listing number of concurrent file accesses, *Vol 2*, 17–7
 - making images header resident, *Vol 2*, 17–18, 17–20
 - making images privileged, *Vol 2*, 17–19, 17–22
 - permanently open images, *Vol 2*, 17–18, 17–20
 - reasons for using, *Vol 2*, 17–17
 - Interactive identifiers, *Vol 1*, 12–11
 - Interactive users
 - limiting for performance management, *Vol 2*, 17–4
 - limiting in system startup, *Vol 1*, 5–16
 - Interchange environment
 - protection, *Vol 1*, 9–27
 - Interfaces to POLYCENTER Software Installation utility
 - DCL, *Vol 1*, 3–14
 - Interfaces to the Backup utility, *Vol 1*, 11–5
 - Internet
 - specifying transport, *Vol 1*, 5–34
 - Intra-Cluster communications, *Vol 1*, 12–9
 - Intrusion databases, *Vol 1*, 12–7
 - Intrusions
 - detection, *Vol 1*, 12–6
 - evasion, *Vol 1*, 12–6
 - Intrusion services
 - managed by security server, *Vol 1*, 5–5
 - IO AUTOCONFIGURE command
 - in SYSMAN, *Vol 1*, 8–7, 8–9
 - in system startup, *Vol 1*, 5–4, 5–7, 8–7
 - IO CONNECT command
 - in SYSMAN, *Vol 1*, 8–9
 - in system startup, *Vol 1*, 5–7
 - IO LOAD command
 - in SYSMAN, *Vol 1*, 8–9
 - IPC (Interrupt Priority C)
 - adjusting quorum, *Vol 1*, 9–78
 - canceling mount verification, *Vol 1*, 9–79
 - description, *Vol 1*, 9–77
 - IPC (Interrupt Priority C) (cont'd)
 - invoking, *Vol 1*, 9–78
 - using to enter XDELTA debugger, *Vol 1*, 9–80
 - IRG (interrecord gap), *Vol 1*, 9–11
 - ISA
 - device configuration, *Vol 1*, 8–17
 - ISL (Initial System Load), *Vol 2*, 25–2
 - ISO 9660 CD
 - file format, *Vol 1*, 9–4
 - ISO 9660 standard
 - data interleaving, *Vol 1*, 9–8
 - duplication of volume label and volume set label, *Vol 1*, 9–50
 - establishing default file attributes, *Vol 1*, 9–31
 - format
 - definition, *Vol 1*, 9–4
 - groups
 - mounting, *Vol 1*, 9–47
 - media
 - protection, *Vol 1*, 9–24
 - showing device information, *Vol 1*, 8–5
 - mounting a volume for, *Vol 1*, 9–30
 - partially mounted volume sets, *Vol 1*, 9–48
 - partially recorded data blocks, *Vol 1*, 9–8
 - restrictions, *Vol 1*, 9–50
 - standard on OpenVMS, *Vol 1*, 9–8
 - UNDEFINED record format errors, *Vol 1*, 9–50
 - using SVDs to mount volumes, *Vol 1*, 9–49
 - volume labels, *Vol 1*, 9–50
 - volume set labels, *Vol 1*, 9–50
 - volume sets
 - mounting, *Vol 1*, 9–47
 - partially mounted, *Vol 1*, 9–48
- ## J
-
- Job banner pages, *Vol 1*, 14–46
 - See also File banner pages
 - \$JOB card, *Vol 1*, 8–39
 - Job controllers
 - See also JOBCTL process; Queue managers and batch jobs, *Vol 1*, 13–2
 - communication with queue manager, *Vol 1*, 13–3
 - relationship to queue manager, *Vol 1*, 13–3, 13–8
 - starting queue manager, *Vol 1*, 13–6, 13–8
 - tasks performed by, *Vol 1*, 13–3
 - JOBCTL process
 - See also Job controllers
 - creation during system startup, *Vol 1*, 5–5
 - Job retention
 - changing for a job, *Vol 1*, 14–29
 - specifying for a job, *Vol 1*, 14–27
 - specifying for a queue, *Vol 1*, 14–27
 - Jobs
 - See also Batch jobs; Output jobs
 - changing scheduling priority, *Vol 1*, 14–36

Jobs (cont'd)

- controlling print position and alignment, *Vol 1*, 14-70, 14-71
 - deleting, *Vol 1*, 14-69
 - holding, *Vol 1*, 14-67, 14-73
 - merging, *Vol 1*, 14-64
 - modifying, *Vol 1*, 14-67
 - moving from one queue to another, *Vol 1*, 14-64
 - releasing, *Vol 1*, 14-67
 - requeuing
 - executing, *Vol 1*, 14-68
 - pending, *Vol 1*, 14-69
 - retaining in a queue, *Vol 1*, 14-29
 - suspending, *Vol 1*, 14-70
- Job scheduling, *Vol 1*, 14-36
- Job status
- definitions, *Vol 1*, 14-66
 - error, *Vol 1*, 14-27, 14-29, 14-34, 14-62
 - holding, *Vol 1*, 14-68, 14-73
 - pending, *Vol 1*, 14-68, 14-69, 14-73
 - retained, *Vol 1*, 14-29, 14-68
 - showing, *Vol 1*, 14-27, 14-65
 - use in job retention, *Vol 1*, 14-27
- Job table quota, *Vol 1*, 7-45
- Journal file of backup information, *Vol 1*, 11-28
- listing contents of, *Vol 1*, 11-28
- Journal file of queue database, *Vol 1*, 13-3
- See also Queue database
 - location, *Vol 1*, 13-6, 13-9
- JTQUOTA process quota, *Vol 1*, 7-45

K

Kernel mode

- calling images running in, *Vol 2*, 17-19, 17-22
 - logical names, *Vol 2*, 17-23
- Keypad definitions, *Vol 2*, 22-6, 22-7
- Key processes
- definition, *Vol 2*, 16-14
- KFMSB adapter
- system tuning, *Vol 2*, 15-37

Kits

See Software products

Known file lists

- definition, *Vol 2*, 17-18
- in system startup, *Vol 1*, 5-13

Known images, *Vol 1*, 10-13

- definition, *Vol 2*, 17-18
- deleting, *Vol 2*, 17-26
- dismounting volume, *Vol 2*, 17-26
- displaying, *Vol 2*, 17-25
- evaluating merits of installing, *Vol 2*, 17-19, 17-25
- file specification for, *Vol 2*, 17-24
- installing, *Vol 2*, 17-24
 - in system startup, *Vol 1*, 5-13; *Vol 2*, 17-18

Known images (cont'd)

- privilege enhancement, *Vol 2*, 17-22
- removing, *Vol 2*, 17-26

L

Labels

- backup tape, *Vol 1*, 11-22
- changing volume, *Vol 1*, 3-30
- header, *Vol 1*, 9-32
- initializing volume to write, *Vol 1*, 9-18
- trailer, *Vol 1*, 9-12

LAD Control Program (LADCP) utility, *Vol 2*, 25-15 to 25-16

- BIND command, *Vol 2*, 25-16
- exiting, *Vol 2*, 25-15
- Help facility, *Vol 2*, 25-16
- invoking, *Vol 2*, 25-15
- making remote InfoServer devices available
 - locally, *Vol 2*, 25-16
- summary, *Vol 2*, 25-15

LADCP (LAD Control Program)

See LAD Control Program (LADCP) utility

LANACP utility

See LAN Auxiliary Control Program utility

LAN Auxiliary Control Program (LANACP) utility

- OPCOM messages displayed, *Vol 2*, 24-27
- running, *Vol 2*, 24-7
- servers, *Vol 2*, 24-6
- stopping, *Vol 2*, 24-7

LAN Control Program (LANCP) utility, *Vol 2*, 24-7

- clearing counters, *Vol 2*, 24-30
- deleting device information, *Vol 2*, 24-22
- deleting node information, *Vol 2*, 24-25
- device database management, *Vol 2*, 24-21
- device management, *Vol 2*, 24-10
- disabling MOP downline load service, *Vol 2*, 24-28
- displaying device information, *Vol 2*, 24-21
- displaying node information, *Vol 2*, 24-23
- displaying OPCOM messages, *Vol 2*, 24-30
- displaying status and counters, *Vol 2*, 24-28, 24-29
- enabling MOP downline load service, *Vol 2*, 24-28
- Load Trace facility, *Vol 2*, 24-30
- MOP console carrier, *Vol 2*, 24-31
- MOP downline load service
 - management, *Vol 2*, 24-28
- MOP trigger boot, *Vol 2*, 24-32
- node database management, *Vol 2*, 24-23
- running, *Vol 2*, 24-7
- setting device information, *Vol 2*, 24-22
- setting node information, *Vol 2*, 24-23
- setting up LAN MOP, *Vol 2*, 24-27
- SPAWN function, *Vol 2*, 24-10

- LAN Control Program (LANCP) utility (cont'd)
 - using command files, *Vol 2*, 24–10
- LANCP utility
 - See LAN Control Program (LANCP) utility
- LAN drivers
 - addresses, *Vol 2*, 24–4
 - multicast, *Vol 2*, 24–4
 - node, *Vol 2*, 24–4
 - physical, *Vol 2*, 24–4
 - Ethernet, *Vol 2*, 24–2
 - FDDI, *Vol 2*, 24–2
 - port, *Vol 2*, 24–2
 - protocol type, *Vol 2*, 24–2
 - Token Ring, *Vol 2*, 24–2
- LAN emulation over ATM
 - device drivers, *Vol 2*, 24–2
- Languages
 - formats, *Vol 1*, 6–17
 - specifying, *Vol 1*, 6–18, 6–19, 6–23
- LANs (local area networks)
 - clearing counters, *Vol 2*, 24–30
 - deleting device information, *Vol 2*, 24–22
 - deleting node information, *Vol 2*, 24–25
 - device management, *Vol 2*, 24–10
 - disabling MOP downline load service, *Vol 2*, 24–28
 - displaying device information, *Vol 2*, 24–21
 - displaying LAN device configurations, *Vol 2*, 24–10
 - displaying LAN device parameters, *Vol 2*, 24–11
 - displaying node information, *Vol 2*, 24–23
 - displaying OPCOM messages, *Vol 2*, 24–30
 - displaying status and counters, *Vol 2*, 24–28, 24–29
 - enabling MOP downline load service, *Vol 2*, 24–28
 - LANACP related OPCOM messages, *Vol 2*, 24–27
 - LAN Auxiliary Control Program (LANACP) utility, *Vol 2*, 24–6
 - LAN Control Program (LANCP) utility, *Vol 2*, 24–7
 - LANCP command files, *Vol 2*, 24–10
 - LANCP SPAWN function, *Vol 2*, 24–10
 - LAN MOP and DECnet MOP, *Vol 2*, 24–25
 - Load Trace facility, *Vol 2*, 24–30
 - migrating DECnet MOP to LAN MOP, *Vol 2*, 24–26
 - MOP console carrier, *Vol 2*, 24–31
 - MOP downline load service, *Vol 2*, 24–25
 - MOP trigger boot, *Vol 2*, 24–32
 - node database management, *Vol 2*, 24–23
 - sample LAN MOP setup, *Vol 2*, 24–27
 - setting device information, *Vol 2*, 24–22
 - setting LAN device parameters, *Vol 2*, 24–13
 - setting node information, *Vol 2*, 24–23
- LANs (local area networks) (cont'd)
 - setting up LAN MOP with CLUSTER_CONFIG, *Vol 2*, 24–26
 - system management enhancements, *Vol 2*, 24–1
 - using multiple LAN adapters for LAT node, *Vol 2*, 26–8
- LASTCP (LASTport Control Program)
 - See LASTport Control Program utility
- LASTport/Disk
 - protocol, *Vol 2*, 25–5
 - service, *Vol 2*, 25–14
- LASTport/Disk service
 - ESS\$DADDRIVER, *Vol 2*, 25–15
- LASTport/Tape
 - protocol, *Vol 2*, 25–5
 - service, *Vol 2*, 25–14
- LASTport/Tape service
 - ESS\$MADDRIVER, *Vol 2*, 25–15
- LASTport Control Program (LASTCP) utility, *Vol 2*, 25–5 to 25–15
 - account requirements, *Vol 2*, 25–14
 - command summary, *Vol 2*, 25–10
 - exiting, *Vol 2*, 25–10
 - functions, *Vol 2*, 25–9, 25–14
 - Help facility, *Vol 2*, 25–10
 - invoking, *Vol 2*, 25–10
 - MAXBUF system parameter requirement, *Vol 2*, 25–14
 - privileges required, *Vol 2*, 25–10
 - transport, *Vol 2*, 25–15
- LAT\$CONFIG.COM command procedure, *Vol 2*, 26–16
 - invoking during system startup, *Vol 1*, 5–15
- LAT\$STARTUP.COM command procedure, *Vol 2*, 26–16, 26–17
 - invoking during system startup, *Vol 1*, 5–15
- LAT\$SYSTARTUP.COM command procedure, *Vol 2*, 26–16, 26–17, 26–18, 26–22
 - example, *Vol 2*, 26–23
 - invoking during system startup, *Vol 1*, 5–15
- LATACP
 - See LAT Ancillary Control Process
- LAT Ancillary Control Process (LATACP), *Vol 2*, 26–24
- LAT Control Program (LATCP) utility, *Vol 2*, 26–3, 26–14
 - See also LAT software
 - DELETE QUEUE_ENTRY command, *Vol 2*, 26–21
 - exiting, *Vol 2*, 26–15
 - features, *Vol 2*, 26–14
 - invoking, *Vol 2*, 26–15
 - /[NO]ANNOUNCEMENTS qualifier, *Vol 2*, 26–4
 - /[NO]LARGE_BUFFER qualifier, *Vol 2*, 26–13
 - queuing incoming requests, *Vol 2*, 26–21

- LAT Control Program (LATCP) utility (cont'd)
 - SET NODE command, *Vol 2, 26-21*
 - SET SERVICE command, *Vol 2, 26-21*
 - setting up limited services, *Vol 2, 26-21*
 - SHOW QUEUE_ENTRY command, *Vol 2, 26-21*
 - SHOW SERVICE command, *Vol 2, 26-21*
 - summary of commands, *Vol 2, 26-15*
 - LATCP (Local Area Transport Control Program)
 - See LAT Control Program (LATCP) utility
 - LAT server, *Vol 2, 25-6*
 - LAT software
 - See also LAT Control Program (LATCP) utility
 - advantages and uses, *Vol 2, 26-2*
 - application programs, *Vol 2, 26-2*
 - creating a service, *Vol 2, 26-19*
 - customizing, *Vol 2, 26-18*
 - disabling service announcements, *Vol 2, 26-4*
 - enabling outgoing connections, *Vol 2, 26-22*
 - load balancing, *Vol 2, 26-3*
 - managing the database size, *Vol 2, 26-24*
 - modems, *Vol 2, 26-2*
 - outgoing connections, *Vol 2, 26-2, 26-3, 26-14*
 - printers, *Vol 2, 26-2*
 - autostart queues on, *Vol 1, 14-4, 14-9*
 - increasing availability of, *Vol 1, 14-4, 14-9*
 - LATSYM symbiont, *Vol 1, 14-3, 14-73*
 - PRTSMB symbiont, *Vol 1, 14-73*
 - sample configuration, *Vol 1, 14-9*
 - setting up, *Vol 1, 14-14*
 - spooling, *Vol 1, 8-36*
 - troubleshooting, *Vol 1, 14-73*
 - queuing incoming requests, *Vol 2, 26-21*
 - service
 - announcements, *Vol 2, 26-4, 26-14*
 - database, *Vol 2, 26-14*
 - dedicated applications, *Vol 2, 26-2*
 - definition, *Vol 2, 26-2*
 - node, *Vol 2, 26-3, 26-4, 26-14*
 - remote printing, *Vol 2, 26-2*
 - setting up limited services, *Vol 2, 26-21*
 - setting up ports, *Vol 2, 26-20*
 - starting network in command procedure, *Vol 1, 5-15; Vol 2, 26-17*
 - starting with LAT\$STARTUP.COM, *Vol 1, 5-15; Vol 2, 26-16, 26-17*
 - terminals, *Vol 1, 8-24; Vol 2, 26-2*
 - determining characteristics of, *Vol 1, 8-24*
 - disconnecting, *Vol 1, 8-23*
 - using large buffers, *Vol 2, 26-13*
 - using multiple LAN adapters, *Vol 2, 26-8*
- LATSYM symbiont, *Vol 1, 14-3, 14-73*
- Layered products
 - startup database, *Vol 1, 5-19*
 - startup phases, *Vol 1, 5-18*
- Layered software installation, *Vol 1, 3-14*
- Lexical functions
 - F\$GETJPI, *Vol 2, 28-9*
 - F\$GETQUI, *Vol 1, 14-59*
 - F\$GETSYI, *Vol 2, 28-9*
 - getting information about queues, *Vol 1, 14-59*
 - getting information about vector processing, *Vol 2, 28-9*
- LIBDECOMP.COM command procedure, *Vol 2, 17-6, 17-9*
- libraries affected, *Vol 2, 17-9*
- parameters, *Vol 2, 17-11*
- using in batch mode, *Vol 2, 17-16*
- using interactively, *Vol 2, 17-11*
- Libraries
 - See Device control libraries
- LIBRARY command
 - used to expand or reduce libraries, *Vol 2, 17-16*
- License database, *Vol 1, 3-5*
 - logical name defining location, *Vol 1, 5-8*
- LICENSE MODIFY command, *Vol 1, 3-5*
- Licenses
 - loading, *Vol 1, 3-4*
 - in system startup, *Vol 1, 5-5*
- Limits
 - See Process limits; UAFs (user authorization files), resource limits
- Line printers
 - preparing for UETP, *Vol 2, 19-3, 19-5, 19-8*
 - testing with UETP, *Vol 2, 19-32*
 - UETP output, *Vol 2, 19-34*
 - UETP test image, *Vol 2, 19-35*
- Lines
 - overflow
 - controlling, *Vol 1, 14-51*
 - TCP/IP
 - connecting hosts, *Vol 2, 23-5*
- Linkable images, *Vol 2, 17-17*
- LINK command, *Vol 1, 10-13*
 - See also Linker utility
 - /SELECTIVE_SEARCH qualifier, *Vol 1, 5-23*
- Linker utility
 - linking against SYS.STB, *Vol 1, 5-23*
- Links
 - hard, *Vol 1, 10-47*
 - primary, *Vol 1, 10-47*
- List operations
 - with BACKUP, *Vol 1, 11-20*
- LM\$JOURNAL file type, *Vol 2, 27-3*
- LMF\$LICENSE logical name
 - defining during system startup, *Vol 1, 5-8*
- LMF (License Management Facility)
 - description of, *Vol 1, 3-4*
 - starting, *Vol 1, 5-5*
- Load balancing
 - using LAT software, *Vol 2, 26-3*

- LOAD command
 - in SYSGEN (VAX), *Vol 1*, 8–8
- Loading device drivers
 - automatically, *Vol 1*, 8–7, 8–9
 - manually, *Vol 1*, 8–9
 - on VAX, *Vol 1*, 8–8
- Load leveling
 - dynamic, *Vol 2*, 28–2
- LOADS logical name, *Vol 2*, 19–35
- Load tests
 - defining user load for UETP, *Vol 2*, 19–14
 - description, *Vol 2*, 19–35
 - error during UETP, *Vol 2*, 19–27
 - running individually, *Vol 2*, 19–13
- Load Trace facility, *Vol 2*, 24–30
- Local Area Network Auxiliary Control Program
 - See LAN Auxiliary Control Program (LANACP) utility
- Local area networks
 - See LANs
- Local identifiers, *Vol 1*, 12–11
- Local page and swap files
 - installing with SATELLITE_PAGE.COM procedure, *Vol 1*, 5–6; *Vol 2*, 16–26
- Location
 - of page file, *Vol 2*, 16–4, 16–6
 - specifying alternate, *Vol 2*, 16–31
 - of queue database, *Vol 1*, 13–4, 13–9
 - master file, *Vol 1*, 13–5
 - queue and journal files, *Vol 1*, 13–6
 - changing, *Vol 1*, 13–6
 - of swap file, *Vol 2*, 16–6
 - specifying alternate, *Vol 2*, 16–31
 - of system dump file, *Vol 2*, 16–2
 - of system files
 - redefining with logical names, *Vol 2*, 17–8
- Log files
 - generated by UETP
 - during the load test, *Vol 2*, 19–28
 - NETSERVER.LOG, *Vol 2*, 19–27
 - OLDUETP.LOG, *Vol 2*, 19–20, 19–21
 - moving to reduce system disk I/O, *Vol 2*, 17–7
 - operator
 - creating new, *Vol 2*, 20–23
 - enabling and disabling classes, *Vol 2*, 20–24
 - maintaining, *Vol 2*, 20–25
 - printing, *Vol 2*, 20–25
 - restarting, *Vol 2*, 20–25
 - security alarm messages, *Vol 2*, 20–22
 - setting up, *Vol 2*, 20–23
 - specifying location, *Vol 2*, 20–23
 - troubleshooting the queue manager, *Vol 1*, 13–16
 - security audit, *Vol 2*, 20–2
 - creating new version, *Vol 2*, 20–31
 - reviewing, *Vol 2*, 20–27
- Logging in, *Vol 1*, 2–10
 - See also Logins
 - when errors in login procedures prevent, *Vol 1*, 4–10
 - when errors in startup procedures prevent, *Vol 1*, 4–10
 - when forgotten passwords prevent, *Vol 1*, 4–11
- Logging out
 - using command procedure, *Vol 1*, 7–22
 - while using VMSINSTAL.COM command procedure, *Vol 1*, 3–3
- Logging startup
 - with SYSMAN, *Vol 1*, 4–17
- Logical disk, for creating a CD-ROM, *Vol 1*, 10–42
- Logical names
 - See also Symbols
 - access modes, *Vol 1*, 5–9
 - ACCOUNTNG, *Vol 2*, 21–4
 - AGEN\$FEEDBACK_REQ_TIME, *Vol 2*, 15–22
 - assigning clusterwide in system startup, *Vol 1*, 5–8
 - assigning for shareable images, *Vol 2*, 17–25
 - assigning systemwide in system startup, *Vol 1*, 5–8
 - assigning to customize the SHUTDOWN.COM procedure, *Vol 1*, 4–28
 - assigning to devices, *Vol 1*, 5–11
 - for software product installations, *Vol 1*, 3–25
 - for system components
 - recommended access mode, *Vol 1*, 5–9
 - LMF\$LICENSE, *Vol 1*, 5–8
 - MAIL\$SYSTEM_FLAGS, *Vol 1*, 5–33
 - MON\$ARCHIVE, *Vol 2*, 20–41
 - NETNODE_REMOTE, *Vol 1*, 5–8
 - NETPROXY, *Vol 1*, 5–8
 - OPC\$LOGFILE_CLASSES, *Vol 2*, 20–24
 - OPC\$LOGFILE_ENABLE, *Vol 2*, 20–24
 - OPC\$LOGFILE_NAME, *Vol 2*, 20–23, 20–25
 - OPC\$OPA0_ENABLE, *Vol 2*, 20–25
 - overriding installation source and destination locations, *Vol 1*, 3–25
 - privilege modes, *Vol 2*, 17–23
 - QMAN\$MASTER, *Vol 1*, 13–5
 - redefining location of system files, *Vol 2*, 17–8
 - RIGHTSLIST, *Vol 1*, 5–8
 - SHOW_CLUSTER\$INIT, *Vol 2*, 22–14
 - SHUTDOWN\$DECNET_MINUTES, *Vol 1*, 4–28
 - SHUTDOWN\$DISABLE_AUTOSTART, *Vol 1*, 4–28, 14–64
 - SHUTDOWN\$INFORM_NODES, *Vol 1*, 4–28, 4–29
 - SHUTDOWN\$MINIMUM_MINUTES, *Vol 1*, 4–28
 - SHUTDOWN\$QUEUE_MINUTES, *Vol 1*, 4–29
 - SHUTDOWN\$TIME, *Vol 1*, 4–29
 - SHUTDOWN\$VERBOSE, *Vol 1*, 4–29

Logical names (cont'd)

- STARTUP\$STARTUP_LAYERED, *Vol 1*, 5–19
 - STARTUP\$STARTUP_VMS, *Vol 1*, 5–19
 - SYS\$ANNOUNCE, *Vol 1*, 5–14
 - SYS\$AUDIT_SERVER_INHIBIT, *Vol 1*, 5–10;
Vol 2, 20–29
 - SYS\$DECDTM_INHIBIT, *Vol 2*, 27–23
 - SYS\$ERRORLOG, *Vol 2*, 20–2
 - SYS\$JOURNAL, *Vol 2*, 27–3
 - SYS\$MONITOR, *Vol 2*, 20–41
 - SYS\$STARTUP, *Vol 1*, 5–2
 - SYS\$SYLOGIN, *Vol 1*, 5–17
 - SYS\$WELCOME, *Vol 1*, 5–15
 - SYSUAF, *Vol 1*, 4–11, 5–8
 - trusted, *Vol 1*, 5–9; *Vol 2*, 17–23
 - used by UETP
 - CTRLNAME, *Vol 2*, 19–33
 - LOADS, *Vol 2*, 19–35
 - SYS\$INPUT, *Vol 2*, 19–32
 - SYS\$OUTPUT, *Vol 2*, 19–34
 - using in SYSMAN, *Vol 1*, 2–15
 - VMSMAIL_PROFILE, *Vol 1*, 5–8
- ## Logical name tables
- definition, *Vol 1*, 5–8
- ## Logical queues
- assigning, *Vol 1*, 14–64
 - description, *Vol 1*, 14–4
 - recommended use, *Vol 1*, 14–4, 14–64
- ## Login command procedures
- booting without, *Vol 1*, 4–10
 - defining announcements in, *Vol 1*, 5–14
 - defining location of during system startup, *Vol 1*, 5–17
 - definition, *Vol 1*, 5–17
 - for captive account, *Vol 1*, 7–32
 - sample, *Vol 1*, 7–33
 - for SYSTEM account, *Vol 1*, 5–17
 - individual, *Vol 1*, 7–19, 7–20
 - sample, *Vol 1*, 7–21
 - in SYSMAN, *Vol 1*, 2–17
 - LOGIN.COM, *Vol 1*, 5–17
 - SYLOGIN.COM, *Vol 1*, 5–17
 - systemwide, *Vol 1*, 5–17, 7–20
 - sample, *Vol 1*, 7–21
 - user-specified, *Vol 1*, 7–20
 - when errors prevent you from logging in, *Vol 1*, 4–10
- ## Logins
- See also Logging in
 - controlling number of dialup attempts, *Vol 1*, 12–6
 - restricting by function, *Vol 1*, 7–30
 - restricting by time, *Vol 1*, 7–28
 - sequence of events, *Vol 1*, 7–5
- ## LOGOUT command, *Vol 2*, 28–9
- ## Logout command procedures
- LOGOUT command, *Vol 1*, 7–23
 - SYLOGOUT.COM, *Vol 1*, 7–22

Long report format

See Console report during UETP

Lost files

recovering, *Vol 1*, 9–71 to 9–72

LPBEGIN phase of system startup, *Vol 1*, 5–19

LPBETA phase of system startup, *Vol 1*, 5–19

LPMAIN phase of system startup, *Vol 1*, 5–19

LTAn devices, *Vol 1*, 8–23

LTDRIVER (LAT port driver)

turning on and off, *Vol 2*, 26–14

M

Machine check errors

if returned when booting, *Vol 1*, 4–18

MAD virtual tape unit, *Vol 2*, 25–14

Magnetic tape ancillary control process

See MTACP

MAIL\$INTERNET_TRANSPORT logical, *Vol 1*, 5–34

Mail utility (MAIL)

customizing in UAF record, *Vol 1*, 7–41

forwarding mail, *Vol 1*, 7–41

logical name for, *Vol 1*, 5–8

MAIL\$SYSTEM_FLAGS logical name, *Vol 1*, 5–33

managing accounts, *Vol 1*, 7–41

sending AUTOGEN report with, *Vol 2*, 15–23

transports, *Vol 1*, 5–34

user profile record

modifying, *Vol 1*, 7–41

removing, *Vol 1*, 7–41

Maintenance operation protocol

See MOP

Management environment

clusterwide, *Vol 1*, 2–16

individual nodes, *Vol 1*, 2–14

local and nonlocal environments, *Vol 1*, 2–14

Manager, queue

See Queue managers

Managing a multiprocessing environment, *Vol 2*, 28–3

tasks for, *Vol 2*, 28–3

Managing a vector processing environment

tasks for, *Vol 2*, 28–5

Managing devices

magnetic tapes

tasks for, *Vol 1*, 8–38

printers

setting characteristics, *Vol 1*, 8–35

tasks for, *Vol 1*, 8–34

tasks for, *Vol 1*, 8–1

terminals

setting characteristics, *Vol 1*, 8–21

tasks for, *Vol 1*, 8–21

Managing file system data caches, *Vol 2*, 18–1

Managing installed software, *Vol 1*, 3–31

Managing page, swap, and dump files
tasks for, *Vol 2*, 16–1

Managing performance
See Performance, managing

Managing system parameters
tasks for, *Vol 2*, 15–1

Managing the LAT database size, *Vol 2*, 26–24

Managing the queue manager and queue database,
Vol 1, 13–1

Mapping pointers
resetting for windows, *Vol 1*, 9–31

Marginal vector consumer
detection of, *Vol 2*, 28–8

Margin size
specifying in forms, *Vol 1*, 14–47

Mass storage control protocol
See MSCP

Master command procedure
See UETP.COM procedure

Master file directories
See MFDs

Master file of queue database, *Vol 1*, 13–3
See also Queue databases
location
specifying, *Vol 1*, 13–5
mounting of disk holding, *Vol 1*, 13–6
QMAN\$MASTER logical name, *Vol 1*, 13–5
saving, *Vol 1*, 13–13

MAXACCTJOBS process limit, *Vol 1*, 7–45

MAXDETACH process limit, *Vol 1*, 7–46

Maximum account jobs process limit, *Vol 1*, 7–45

Maximum detached process limit, *Vol 1*, 7–46

MAXJOBS process limit, *Vol 1*, 7–46

MAXSYSGROUP system parameter, *Vol 1*, 12–8

MAX_DUMPFILE symbol, *Vol 2*, 16–31

MAX_ERRLOGDUMP symbol, *Vol 2*, 16–31

MAX_PAGEFILE_n_SIZE symbol, *Vol 2*, 16–31

MAX_PAGEFILE symbol, *Vol 2*, 16–31

MAX_ prefix for AUTOGEN, *Vol 2*, 15–20

MAX_SWAPFILE_n_SIZE symbol, *Vol 2*, 16–31

MAX_SWAPFILE symbol, *Vol 2*, 16–31

Media
definition, *Vol 1*, 9–2

Media errors
analyzing, *Vol 1*, 9–80

Memory
allotted to vector consumer processes, *Vol 2*,
28–7
cache, *Vol 2*, 18–1
conserving with shareable images, *Vol 2*, 17–21
deferred testing, *Vol 1*, 4–3
images in, *Vol 2*, 17–19
information captured in crash dump, *Vol 2*,
16–2
physical dump, *Vol 2*, 16–4, 16–12

Memory
information captured in crash dump (cont'd)
selective dump, *Vol 2*, 16–4, 16–12
making efficient use of by installing images,
Vol 2, 17–7
paging, *Vol 2*, 16–5
sections in, *Vol 2*, 17–21
swapping, *Vol 2*, 16–5
when large sizes prevent storing a complete
system dump, *Vol 2*, 16–12

Memory-resident sections, *Vol 2*, 17–26

MEMORY_TEST environment variable
deferred memory testing, *Vol 1*, 4–3

Messages
defining welcome, *Vol 1*, 5–14

DIBOL
starting the DIBOL message manager, *Vol*
1, 5–16
enabling and disabling, *Vol 2*, 20–18
error
removing, *Vol 2*, 22–10
indicating execution of startup command
procedures
site-independent, *Vol 1*, 4–6
site-specific, *Vol 1*, 4–6
indicating high page or swap file fragmentation,
Vol 2, 16–34
indicating insufficient page file size, *Vol 2*,
16–10
indicating lack of installed page file, *Vol 1*, 5–6
indicating successful boot, *Vol 1*, 4–6
indicating that a vector processor is not
available, *Vol 2*, 28–6
indicating that login is possible, *Vol 1*, 4–6
login welcome, *Vol 1*, 2–10

OPCOM
See OPCOM messages
operator replies, *Vol 1*, 2–26; *Vol 2*, 20–20
operator requests, *Vol 1*, 2–26
question mark (?), *Vol 1*, 4–18
saving in a file, *Vol 1*, 3–26
security alarm, *Vol 2*, 20–22
sending to users with OPCOM, *Vol 1*, 2–23
suppressing the display, *Vol 2*, 15–12
user requests, *Vol 2*, 20–20
using when managing a system, *Vol 1*, 2–5
WRITEBOOT, *Vol 1*, 4–21

MFDs (master file directories)
BACKUP stores save-set file in, *Vol 2*, A–8
contains directory structure for volume set, *Vol*
1, 9–46
definition, *Vol 2*, A–8
on root volume in volume set, *Vol 1*, 9–44
reserved file, *Vol 2*, A–8
reserved files listed in, *Vol 1*, 9–5

MIB subagents
implemented on OpenVMS Alpha, *Vol 2*, C–1

MIME utility, *Vol 1*, 5–34
 setting up, *Vol 1*, 5–34

Minimerge and volume shadowing, *Vol 1*, 11–45

Minimum startup
 booting with, *Vol 1*, 4–15

MIN_DUMPFILE symbol, *Vol 2*, 16–31

MIN_ERRLOGDUMP symbol, *Vol 2*, 16–31

MIN_PAGEFILE n _SIZE symbol, *Vol 2*, 16–31

MIN_PAGEFILE symbol, *Vol 2*, 16–31

MIN_ prefix for AUTOGEN, *Vol 2*, 15–20

MIN_SWAPFILE n _SIZE symbol, *Vol 2*, 16–31

MIN_SWAPFILE symbol, *Vol 2*, 16–31

MMG_CTLFLAGS system parameter
 deferred memory testing, *Vol 1*, 4–3

MODE logical name, *Vol 2*, 19–18, 19–39

Modes
 See Executive mode; Privilege mode

Modifiable startup command procedure
 See Site-specific startup command procedure

Modifying size of page, swap, and dump files
 See Changing size of page, swap, and dump files

Modifying system parameters
 See Changing system parameters

MODPARAMS.DAT file, *Vol 2*, 15–17, 15–18

ADD_ prefix, *Vol 2*, 15–19

controlling page, swap, and dump files, *Vol 2*, 16–29
 specifying location, *Vol 2*, 16–31
 specifying size of individual files, *Vol 2*, 16–31
 specifying size of total file space, *Vol 2*, 16–31

controlling parameter values set by AUTOGEN, *Vol 2*, 15–4, 15–18

creating page, swap, and dump files, *Vol 2*, 16–31

including external parameter files in, *Vol 2*, 15–22

increasing parameter values, *Vol 2*, 15–19

logging statements, *Vol 2*, 15–12

MAX_ prefix, *Vol 2*, 15–20

MIN_ prefix, *Vol 2*, 15–20

sample, *Vol 2*, 15–17

specifying an alternate default startup
 command, *Vol 1*, 4–15

specifying parameter values
 absolute, *Vol 2*, 15–20
 maximum, *Vol 2*, 15–20
 minimum, *Vol 2*, 15–20

storing your system parameter changes in, *Vol 2*, 15–5

Modules
 device control
 See Device control modules

MONITOR.COM command procedure, *Vol 2*, 20–41
 used with Monitor utility, *Vol 2*, 20–41

MONITOR command
 See also Monitor utility
 controlling amount of time between displays of
 screen images, *Vol 2*, 20–38
 recording data on time spent, *Vol 2*, 20–37
 recording file system and process data, *Vol 2*, 20–38
 routing display output, *Vol 2*, 20–36
 specifying how log request is to run, *Vol 2*, 20–35
 specifying input file, *Vol 2*, 20–38
 specifying node, *Vol 2*, 20–37
 specifying time of display, *Vol 2*, 20–38
 storing display output, *Vol 2*, 20–39

Monitoring
 multiprocessing environment, *Vol 2*, 28–3

Monitor utility (MONITOR)
 See also MONITOR command
 analyzing disk use with, *Vol 1*, 9–14
 class types, *Vol 2*, 20–32
 command procedures, *Vol 2*, 20–41
 for cluster summaries, *Vol 2*, 20–41
 to initiate continuous recording, *Vol 2*, 20–43
 to produce cluster summaries, *Vol 2*, 20–43
 description, *Vol 2*, 20–32
 directing display, *Vol 2*, 20–34
 displaying and recording concurrently, *Vol 2*, 20–37
 displaying live monitoring, *Vol 2*, 20–35
 entering commands, *Vol 2*, 20–35
 exiting, *Vol 2*, 20–35
 invoking, *Vol 2*, 20–35
 logical name for, *Vol 1*, 5–8

MONITOR.COM command procedure
 using to create summary file, *Vol 2*, 20–41

MONSUM.COM command procedure
 using to generate clusterwide multfile
 summary reports, *Vol 2*, 20–41

moving log file to reduce system disk I/O, *Vol 2*, 17–7

parameters, *Vol 2*, 20–33

playing back monitoring, *Vol 2*, 20–38

playing back remote monitoring, *Vol 2*, 20–40

producing reports, *Vol 2*, 20–41

qualifiers, *Vol 2*, 20–34

recording live monitoring, *Vol 2*, 20–37

reports, *Vol 2*, 20–41

rerecording monitoring, *Vol 2*, 20–40

running continuously, *Vol 2*, 20–41

SUBMON.COM procedure
 using to start MONITOR.COM as a
 detached process, *Vol 2*, 20–41

version compatibility, *Vol 2*, 20–45

MONSUM.COM command procedure, *Vol 2*, 20–43
 used with Monitor utility, *Vol 2*, 20–41
 using to generate clusterwide multifile summary reports, *Vol 2*, 20–41

MOP (maintenance operation protocol)
 clearing counters, *Vol 2*, 24–30
 coexisting with DECnet MOP, *Vol 2*, 24–25
 console carrier, *Vol 2*, 24–31
 disabling, *Vol 2*, 24–28
 displaying status and counters, *Vol 2*, 24–28, 24–29
 downline loading, *Vol 2*, 25–2
 enabling, *Vol 2*, 24–28
 functions of, *Vol 2*, 23–11
 LANACP server, *Vol 2*, 24–28
 LAN MOP, *Vol 2*, 24–25
 migrating DECnet MOP to LAN MOP, *Vol 2*, 24–26
 sample setup of LAN MOP, *Vol 2*, 24–27
 setup with CLUSTER_CONFIG.COM, *Vol 2*, 24–26
 trigger boot, *Vol 2*, 24–32

MOUNT command
 See also Mounting volumes
 adding to disk volume sets, *Vol 1*, 9–29
 assigning a volume set name, *Vol 1*, 9–44
 avoiding use of /CLUSTER with SYSMAN DO command, *Vol 2*, 22–18
 controlling whether header labels are written to a volume, *Vol 1*, 9–32
 copying files to tape, *Vol 1*, 10–38
 creating a public volume, *Vol 1*, 9–31
 creating disk volume sets, *Vol 1*, 9–29, 9–43
 disabling automatic notification of mount failures, *Vol 1*, 9–29
 disabling automatic volume switching, *Vol 1*, 9–54
 disabling mount verification feature for disks, *Vol 1*, 9–30
 disabling mount verification feature for tapes, *Vol 1*, 9–32
 disk and tape volumes, *Vol 1*, 9–27
 enabling automatic notification of mount failures, *Vol 1*, 9–29
 enabling mount verification feature for disks, *Vol 1*, 9–30
 enabling mount verification feature for tapes, *Vol 1*, 9–32
 enabling the processing of Subsystem Aces, *Vol 1*, 9–31
 enabling the write cache for a tape device, *Vol 1*, 9–32
 ensuring that tape volume set has been initialized, *Vol 1*, 9–54
 establishing default file attributes for records on ISO 9660 media, *Vol 1*, 9–31
 foreign volumes, *Vol 1*, 9–30, 9–32, 10–28

MOUNT command (cont'd)
 including a quoted text string as part of mount request, *Vol 1*, 9–30
 inhibiting access checks, *Vol 1*, 9–32
 in OpenVMS Cluster environment, *Vol 1*, 9–28, 9–30
 in system startup
 for remote InfoServer devices, *Vol 2*, 25–16
 mounting page and swap file disks, *Vol 1*, 5–6
 special consideration about operator assistance, *Vol 1*, 5–12
 ISO 9660 media, *Vol 1*, 9–30
 overriding expiration date, *Vol 1*, 10–29
 overriding protection, *Vol 1*, 9–33
 overriding protection checks, *Vol 1*, 9–30
 overriding the volume identification field, *Vol 1*, 9–32
 overriding UIC in second volume label, *Vol 1*, 9–32
 parameters, *Vol 1*, 9–28
 protection codes, *Vol 1*, 10–28
 public volumes, *Vol 1*, 9–28
 qualifiers, *Vol 1*, 9–29, 9–31
 requesting operator assistance, *Vol 1*, 9–29
 resetting the number mapping pointers, *Vol 1*, 9–31
 specifying block size for tape, *Vol 1*, 9–32
 specifying number of bytes in each record, *Vol 1*, 9–33
 specifying record size, *Vol 1*, 9–33
 specifying that other users can access current volume, *Vol 1*, 9–31
 specifying the number of directories that the system keeps in memory, *Vol 1*, 9–29
 specifying the number of disk blocks allocated, *Vol 1*, 9–30
 specifying UIC, *Vol 1*, 9–32
 suspending quota operation on a volume, *Vol 1*, 9–67
 to mount disk holding page and swap files, *Vol 2*, 16–26

Mounting forms, *Vol 1*, 14–51
 matching stock, *Vol 1*, 14–47

Mounting of queue database disk, *Vol 1*, 13–6

Mounting volumes, *Vol 1*, 11–16
 See also MOUNT command
 disks, *Vol 1*, 9–27
 for queue database files, *Vol 1*, 13–6
 in system startup, *Vol 1*, 5–11
 early, *Vol 1*, 5–12
 for page and swap files, *Vol 1*, 5–6
 InfoServer, *Vol 2*, 25–16
 special consideration about operator assistance, *Vol 1*, 5–12
 virtual device unit, *Vol 2*, 25–16
 if device is unavailable, *Vol 1*, 9–35

Mounting volumes (cont'd)

- in an OpenVMS Cluster environment, *Vol 1*, 9-28
- operator assistance, *Vol 1*, 9-25
- public, *Vol 1*, 9-28
- substituting, *Vol 1*, 9-35
- tape, *Vol 1*, 9-27
- tape volume sets, *Vol 1*, 9-51
- with caching disabled, *Vol 2*, 18-3

Mounting volume sets

See also MOUNT command

- disk, *Vol 1*, 9-44, 9-45, 9-46
- tape
 - with automatic volume switching disabled, *Vol 1*, 9-54

Mount messages

- disabling with SUBSYSTEM qualifier, *Vol 1*, 9-31

Mount utility (MOUNT)

- using to mount ISO 9660 volume sets, *Vol 1*, 9-48

Mount verification

- aborted
 - OPCOM message, *Vol 1*, 9-79
- aborting by dismounting, *Vol 1*, 9-77
- canceling, *Vol 1*, 9-77
 - using IPC, *Vol 1*, 9-79
- definition, *Vol 1*, 9-73
- device
 - off line, *Vol 1*, 9-74, 9-75
 - write-locked, *Vol 1*, 9-76
- disabling, *Vol 1*, 9-30
 - for tapes, *Vol 1*, 9-32
- enabling, *Vol 1*, 9-30, 9-75
 - for tapes, *Vol 1*, 9-32
- messages, *Vol 1*, 9-75
- operation of, *Vol 1*, 9-73
- timeout, *Vol 1*, 9-75
 - OPCOM message, *Vol 1*, 9-75

MOVE keypad function, *Vol 2*, 22-7

Moving jobs from one queue to another, *Vol 1*, 14-64

MSCP (mass storage control protocol), *Vol 2*, 22-2

MSGHLP

- See Help Message utility

.MSGHLP\$DATA files

- adding to the Help Message database, *Vol 1*, 5-28

MSGHLP\$LIBRARY.MSGHLP\$DATA file, *Vol 1*, 5-28

MSGHLP\$LIBRARY logical name, *Vol 1*, 5-28

MTACP (magnetic tape ancillary control process)

- definition, *Vol 1*, 9-11

Multiple queue managers, *Vol 1*, 13-3, 13-10

- commands affected by, *Vol 1*, 13-11
- managing queues with, *Vol 1*, 13-12
- moving queues, *Vol 1*, 13-12
- restriction, *Vol 1*, 13-3

Multiple queue managers (cont'd)

- specifying names of, *Vol 1*, 13-11
- specifying queue manager name, *Vol 1*, 13-11, 13-12
- use of queue database, *Vol 1*, 13-11

Multiple tape densities

- using on Alpha systems, *Vol 1*, 9-12

Multiprocessing, *Vol 2*, 28-2

- definition, *Vol 2*, 28-2
- displaying information, *Vol 2*, 28-3
- monitoring, *Vol 2*, 28-3
- system parameters, *Vol 2*, 28-3

MULTIPROCESSING system parameter, *Vol 2*, 28-3

MVTIMEOUT system parameter, *Vol 1*, 9-75, 9-77

N

Namespaces

- in DECnet-Plus full name, *Vol 2*, 23-10

Naming conventions

- for queue and journal files for additional queue managers, *Vol 1*, 13-4
- of devices, *Vol 1*, 8-1
 - in an OpenVMS Cluster environment, *Vol 1*, 8-1
 - virtual terminals, *Vol 1*, 8-22

NCL (Network Control Language), *Vol 2*, 23-9

NCP (Network Control Program)

- controlling proxy access, *Vol 1*, 7-40
- for managing Phase IV nodes, *Vol 2*, 23-10

NET\$CONFIGURE

- configuring DECnet-Plus, *Vol 2*, 23-11

NET\$PROXY.DAT files, *Vol 1*, 7-4, 7-37

- moving to reduce system disk I/O, *Vol 2*, 17-7
- use of, *Vol 1*, 12-7

NET\$PROXY logical name

- defining to reduce system disk I/O, *Vol 2*, 17-7

NETDRIVER (network driver)

- connecting, *Vol 1*, 8-9, 8-10

NETNODE_REMOTE logical name

- defining during system startup, *Vol 1*, 5-8

NETPROXY.DAT files, *Vol 1*, 7-4, 7-37

- moving to reduce system disk I/O, *Vol 2*, 17-7

NETPROXY logical name

- defining during system startup, *Vol 1*, 5-8
- defining to reduce system disk I/O, *Vol 2*, 17-7

Network communications device

- connecting, *Vol 1*, 8-9, 8-10

Network Control Language

- See NCL

Network identifiers, *Vol 1*, 12-11

Network management

- concepts, *Vol 2*, 23-5, 23-10
- distributed, *Vol 2*, 23-5, 23-10

Network proxy accounts, *Vol 1*, 7–37

Networks

- communications device, *Vol 1*, 8–10
- Compaq networking options, *Vol 2*, 23–2
- LAN enhancements, *Vol 2*, 24–1
- networking software options, *Vol 2*, 23–2
- possible software combinations, *Vol 2*, 23–3
- proxy database
 - creating, *Vol 1*, 7–37
 - logical name defining location, *Vol 1*, 5–8
- remote node database
 - logical name defining location, *Vol 1*, 5–8
- save sets, *Vol 1*, 11–8
- starting in system startup, *Vol 1*, 5–16

Node names

- DECnet Phase IV, *Vol 2*, 23–9
- DECnet-Plus full names, *Vol 2*, 23–9
- naming conventions, *Vol 2*, 23–9
- required for OpenVMS InfoServer Client startup, *Vol 2*, 25–11

Nodes

- deleting node information, *Vol 2*, 24–25
- in LAT database, *Vol 2*, 26–14
- local
 - definition, *Vol 1*, 2–13
 - setting node information, *Vol 2*, 24–23
 - transferring files between, *Vol 1*, 10–40

Nondeductible resource, *Vol 1*, 7–3

Nonstop boot

- definition, *Vol 1*, 4–4
- performing, *Vol 1*, 4–4

NPAGEDYN system parameter, *Vol 2*, 28–7

NUM_ETHERADAPT symbol, *Vol 2*, 15–21

NUM_NODES symbol, *Vol 2*, 15–21

O

Objects

- protecting volume, *Vol 1*, 9–21

ODS (On-Disk Structure)

- CD on OpenVMS, *Vol 1*, 9–8
- comparing levels, *Vol 1*, 9–6
- definition, *Vol 1*, 9–2, 9–3
- disk and CD-ROM file structures, *Vol 1*, 9–4
- Levels 1, 2, and 5, *Vol 1*, 9–4
- on dual format CDs, *Vol 1*, 9–4
- reserved files, *Vol 1*, 9–5

ODS-1

- comparing with ODS-2 and ODS-5, *Vol 1*, 9–6
- migrating to ODS-2, *Vol 1*, 9–39

ODS-2

- comparing with ODS-1, *Vol 1*, 9–6
- converting from ODS-1, *Vol 1*, 9–39
- converting from ODS-5, *Vol 1*, 9–39
- converting to ODS-5, *Vol 1*, 9–36
- displaying type of volume, *Vol 1*, 10–3

ODS-5

- backing up, *Vol 1*, 10–9
- case of file name, *Vol 1*, 10–4
- comparing with ODS-1 and ODS-2, *Vol 1*, 9–6
- controlling access to, *Vol 1*, 10–14
- converting an existing volume to, *Vol 1*, 9–36
- converting to ODS-2, *Vol 1*, 9–39
- copying contents of disk to ODS-2 disk, *Vol 1*, 9–40
- copying files to ODS-2 disk, *Vol 1*, 9–40
- description, *Vol 1*, 10–2
- displaying type of volume, *Vol 1*, 10–3
- extended file names
 - creating, *Vol 1*, 10–3
 - on Alpha systems, *Vol 1*, 10–8
 - on VAX systems, *Vol 1*, 10–8
- initializing a new volume, *Vol 1*, 9–20
- mounting a new volume, *Vol 1*, 9–20
- preventing untested applications from accessing, *Vol 1*, 10–14
- preventing VAX users from accessing, *Vol 1*, 10–14
- restoring save set to ODS-2 disk, *Vol 1*, 9–39
- saving disk to ODS-2 save set, *Vol 1*, 9–39
- saving files in ODS-2 save set, *Vol 1*, 9–41

ODS-5 volumes

- controlling access to, *Vol 1*, 10–14

ODS volumes

- converting from ODS-1 to ODS-2, *Vol 1*, 9–39
- OLDUETP.LOG file, *Vol 2*, 19–20
- OPA0: device, *Vol 1*, 2–24
- OPC\$LOGFILE_CLASSES logical name, *Vol 2*, 20–24
- OPC\$LOGFILE_ENABLE logical name, *Vol 2*, 20–24
- OPC\$LOGFILE_NAME logical name, *Vol 2*, 20–25
 - for operator log file, *Vol 2*, 20–23
- OPC\$OPA0_ENABL logical name, *Vol 2*, 20–25
- OPCCRASH.EXE program, *Vol 1*, 4–22
- OPCOM (Operator Communication Manager), *Vol 1*, 11–18
 - See also OPCOM messages; Operator log file
 - automatic restart, *Vol 1*, 2–23
 - classes
 - enabling, *Vol 1*, 2–24
 - enabling and disabling for log file, *Vol 2*, 20–24
 - communicating with operators, *Vol 1*, 9–25
 - communicating with users, *Vol 1*, 9–25
 - components of, *Vol 1*, 2–20
 - default behavior, *Vol 1*, 2–22
 - disabling operator terminals, *Vol 1*, 2–25
 - enabling operator classes, *Vol 1*, 2–24
 - enabling operator terminals, *Vol 1*, 2–24
 - failure, *Vol 1*, 2–23
 - illustration, *Vol 1*, 2–20
 - log file, *Vol 1*, 2–22
 - messages

- OPCOM (Operator Communication Manager)
 - messages (cont'd)
 - See OPCOM messages
 - mount verification, *Vol 1*, 9–74
 - operator log file
 - creating new version, *Vol 2*, 20–18
 - printing, *Vol 2*, 20–25
 - restarting, *Vol 2*, 20–25
 - setting up, *Vol 2*, 20–23
 - specifying default state, *Vol 2*, 20–24
 - operator terminals, *Vol 1*, 2–22
 - process, *Vol 1*, 2–22
 - creation during system startup, *Vol 1*, 5–5
 - process dump file, *Vol 1*, 2–23
 - replying to operator requests, *Vol 1*, 2–26
 - requirements, *Vol 1*, 2–22
 - sending requests to an operator, *Vol 1*, 2–26
 - startup of, *Vol 1*, 2–23
 - uses of, *Vol 1*, 2–20
 - using to request assistance, *Vol 1*, 11–18
- OPCOM.DMP process dump file, *Vol 1*, 2–23
- OPCOM messages, *Vol 1*, 2–22
 - continuation volume request, *Vol 1*, 9–54, 10–39
 - controlling, *Vol 1*, 2–23
 - displayed during LANACP start, *Vol 2*, 24–27
 - generated by LANACP, *Vol 2*, 24–30
 - mount request message, *Vol 1*, 9–34
 - mount verification, *Vol 1*, 9–75
 - aborted message, *Vol 1*, 9–79
 - timeout message, *Vol 1*, 9–75
 - operator reply, *Vol 2*, 20–20
 - security alarm, *Vol 2*, 20–22
 - sending, *Vol 1*, 2–23
 - SYSGEN, *Vol 2*, 20–21
 - user request, *Vol 2*, 20–20
- Open file limit, *Vol 1*, 7–45
- Open images, *Vol 2*, 17–18, 17–20
- OpenVMS Cluster environments
 - dual-architecture
 - installing images, *Vol 2*, 22–18
 - preparing for UETP, *Vol 2*, 19–11
 - security audit log files in, *Vol 2*, 20–31
 - test failure during UETP, *Vol 2*, 19–26
- OpenVMS Cluster systems
 - adjusting quorum after shutting down a node, *Vol 1*, 4–24
 - autostart queues in, *Vol 1*, 14–4
 - benefits, *Vol 2*, 22–2
 - common parameter files in, *Vol 2*, 15–22
 - cross-architecture booting, *Vol 2*, 24–28
 - DECnet node configuration, *Vol 2*, 23–10
 - defining in SYSMAN, *Vol 1*, 2–16
 - defining location of queue database files, *Vol 1*, 13–5
 - device names, *Vol 1*, 8–1
 - disks in, *Vol 1*, 9–3
 - dismounting a volume, *Vol 1*, 9–60
- OpenVMS Cluster systems (cont'd)
 - dual-architecture
 - installing images, *Vol 2*, 22–18
 - executing commands in, *Vol 1*, 2–11
 - generic batch queues in, *Vol 1*, 14–6
 - generic output queues in, *Vol 1*, 14–11
 - generic queues in, *Vol 1*, 14–3
 - LAN management enhancements, *Vol 2*, 24–1
 - LAN MOP coexisting with DECnet MOP, *Vol 2*, 24–25
 - limiting nodes that can run the queue manager, *Vol 1*, 13–9
 - local and nonlocal, *Vol 1*, 2–16
 - managing with SYSMAN, *Vol 2*, 22–16 to 22–20
 - migrating from DECnet MOP to LAN MOP, *Vol 2*, 24–26
 - monitoring with SHOW CLUSTER, *Vol 2*, 22–4
 - mounting volumes in, *Vol 1*, 9–28
 - moving the queue manager to another node, *Vol 1*, 13–10
 - node restriction for a startup command procedure, *Vol 1*, 5–20
 - print and batch queues in, *Vol 2*, 22–2
 - remote monitoring limitation, *Vol 2*, 20–45
 - security, *Vol 2*, 22–16
 - setting up for LAN MOP, *Vol 2*, 24–27
 - shared resources, *Vol 2*, 22–2
 - shutting down an entire OpenVMS Cluster, *Vol 1*, 4–24
 - specifying location of queue database, *Vol 1*, 13–6
 - specifying preferred order of queue manager nodes, *Vol 1*, 13–9, 13–10
 - support for TCP/IP, *Vol 2*, 23–5
 - SYSMAN management environment, *Vol 1*, 2–13
 - system management, *Vol 2*, 22–3
 - using CLUSTER_CONFIG.COM, *Vol 2*, 22–2
 - using CLUSTER_CONFIG to set up LAN MOP, *Vol 2*, 24–26
- OpenVMS Cluster Systems
 - problems and restrictions
 - tuning BAP system parameters, *Vol 2*, 15–37
- OpenVMS Management Station
 - description, *Vol 1*, 2–2
 - documentation, *Vol 1*, 2–4
 - features, *Vol 1*, 2–4
- OpenVMS Registry
 - backup, *Vol 2*, D–11
 - connecting to a Windows NT system, *Vol 2*, D–9
 - failover in a cluster, *Vol 2*, D–9
 - installing, *Vol 2*, D–1
 - quotas, *Vol 2*, D–10
 - restoring, *Vol 2*, D–11

- OpenVMS Registry (cont'd)
 - running in an OpenVMS Alpha mixed-version cluster, *Vol 2, D-11*
 - security, *Vol 2, D-10*
 - shutting down, *Vol 2, D-6*
 - starting, *Vol 2, D-5*
 - Unicode support, *Vol 2, D-11*
 - Utilities for configuring, *Vol 2, D-1*
- OpenVMS Registry Configuration utility
 - menu, *Vol 2, D-1*
 - options, *Vol 2, D-2*
- OpenVMS Registry server commands, *Vol 2, D-6*
- Operating systems
 - building on another disk, *Vol 1, 2-28*
 - copying to another system disk, *Vol 1, 2-31*
- OPERATOR.LOG file
 - See Operator log file
- Operator assistance
 - operator classes, *Vol 1, 2-24*
 - replying to operator requests, *Vol 1, 2-26*
 - with MOUNT command, *Vol 1, 9-25*
- Operator classes
 - See OPCOM, classes
- Operator communication manager
 - See OPCOM
- Operator consoles
 - enabling in system startup, *Vol 1, 5-5*
- Operator log file
 - See also OPCOM; OPCOM messages
 - access to, *Vol 2, 20-23*
 - closing current and opening new, *Vol 2, 20-18*
 - creating new, *Vol 2, 20-23*
 - definition, *Vol 1, 2-22*
 - device status message, *Vol 2, 20-18, 20-23*
 - enabling and disabling classes, *Vol 2, 20-24*
 - enabling in system startup, *Vol 1, 5-5*
 - initialization message, *Vol 2, 20-18*
 - logging user requests in, *Vol 2, 20-20*
 - maintaining, *Vol 2, 20-25*
 - printing, *Vol 2, 20-23, 20-25*
 - purging, *Vol 2, 20-25*
 - during system startup, *Vol 1, 5-14*
 - recording changes to system parameters, *Vol 2, 20-21*
 - request identification number, *Vol 2, 20-20*
 - response recorded in, *Vol 2, 20-20*
 - restarting, *Vol 2, 20-25*
 - security alarm messages in, *Vol 2, 20-22*
 - setting up, *Vol 2, 20-17, 20-23*
 - size, *Vol 2, 20-23*
 - specifying default state, *Vol 2, 20-24*
 - terminal enable and disable message, *Vol 2, 20-23*
 - troubleshooting the queue manager, *Vol 1, 13-16*
 - user request and operator reply messages, *Vol 2, 20-23*
- Operator log file (cont'd)
 - volume mount and dismount messages, *Vol 2, 20-23*
- Operators (computer)
 - assistance with MOUNT command in system startup, *Vol 1, 5-12*
 - classes of, *Vol 1, 2-24*
 - replying to requests, *Vol 1, 2-26*
 - requesting assistance from, *Vol 1, 9-25*
 - sending requests to, *Vol 1, 2-26*
- Operator terminals, *Vol 1, 2-22*
 - designating, *Vol 1, 2-24*
 - in batch or SYSTARTUP, *Vol 1, 2-22*
 - enabling and disabling, *Vol 1, 2-24, 2-25; Vol 2, 20-18*
 - security alarms, *Vol 2, 20-30*
 - security terminal, *Vol 2, 20-22*
 - setting up, *Vol 1, 9-25*
 - user request, *Vol 1, 9-25*
- Option list
 - parameter for VMSINSTAL.COM command procedure, *Vol 1, 3-8*
- Options
 - for software product installations, *Vol 1, 3-26*
- OSI
 - See Open Systems Interconnection
- OSI (Open Systems Interconnection), *Vol 2, 23-9*
- Output a newly created PCF, *Vol 1, 3-26*
- Output devices
 - See Printers; Terminals
- Output execution queues
 - See also Execution queues; Output queues
 - definition, *Vol 1, 14-3*
- Output jobs
 - See also Output queues
 - aligning forms, *Vol 1, 14-71*
 - allowing to complete before stopping a queue, *Vol 1, 14-63*
 - changing scheduling priority, *Vol 1, 14-36*
 - controlling, *Vol 1, 14-65*
 - controlling print position and alignment, *Vol 1, 14-70, 14-71*
 - deleting, *Vol 1, 14-69*
 - holding and releasing, *Vol 1, 14-67*
 - modifying, *Vol 1, 14-67*
 - monitoring, *Vol 1, 14-65*
 - requeuing
 - executing, *Vol 1, 14-68*
 - pending, *Vol 1, 14-69*
 - resuming printing, *Vol 1, 14-70, 14-71*
 - retaining in a queue, *Vol 1, 14-29*
 - scheduling, *Vol 1, 14-36*
 - status
 - See Job status
 - suspending, *Vol 1, 14-70*

Output queues

- See also Output jobs; Output queuing environment
- assigning a default form, *Vol 1*, 14–50
- canceling characteristics, *Vol 1*, 14–31
- changing DEFAULT form, *Vol 1*, 14–48
- commands for managing, *Vol 1*, 14–57
- creating, *Vol 1*, 14–15
- creating a form, *Vol 1*, 14–49
- deleting, *Vol 1*, 14–65
- execution
 - description, *Vol 1*, 14–3
 - printer, *Vol 1*, 14–3
 - server, *Vol 1*, 14–3
 - terminal, *Vol 1*, 14–3
- mounting a form on, *Vol 1*, 14–51
- options, *Vol 1*, 14–20
 - banner pages, *Vol 1*, 14–38
 - controlling page and line overflow, *Vol 1*, 14–51
 - device control libraries, *Vol 1*, 14–52
 - forms, *Vol 1*, 14–46
 - qualifiers for specifying, *Vol 1*, 14–21 to 14–22
 - restricting access, *Vol 1*, 14–23
 - retaining jobs, *Vol 1*, 14–27
 - suppressing form feed, *Vol 1*, 14–52
- order of device control module output, *Vol 1*, 14–53
- pausing, *Vol 1*, 14–61, 14–70
 - to align position of print for preprinted forms, *Vol 1*, 14–70, 14–71
 - to change position of print, *Vol 1*, 14–70
- rerouting jobs in, *Vol 1*, 14–64
- specifying page and margin size in forms, *Vol 1*, 14–47
- standalone workstations, *Vol 1*, 14–8
- starting, *Vol 1*, 14–15
- status, *Vol 1*, 14–59
- stopping, *Vol 1*, 14–62, 14–63

Output queuing environment

- for LAT printers, *Vol 1*, 14–9
- for mixed printers, *Vol 1*, 14–8
- for multiple printers of the same kind, *Vol 1*, 14–10
- in OpenVMS Cluster environments, *Vol 1*, 14–11
- sample configurations, *Vol 1*, 14–8 to 14–13
- single printer, *Vol 1*, 14–8
- spooled printers, *Vol 1*, 14–12
- standalone workstation, *Vol 1*, 14–8
- steps for setting up, *Vol 1*, 14–13

Overdraft limit

- user exceeding quota, *Vol 1*, 9–64

Overflow

- See Lines, overflow; Page overflow

Overriding

- DECnet parameters, *Vol 2*, 15–21

Owner

- category of user access, *Vol 1*, 12–8

Ownership

- file
 - displaying, *Vol 1*, 10–22

P

P1 space

- IMGREG_PAGES system parameter determines size, *Vol 2*, 17–21
- shared address data assigns, *Vol 2*, 17–21

Packets, *Vol 2*, 23–5

Page and swap files, *Vol 1*, 10–13

PAGEFILE.SYS file, *Vol 2*, 16–3, 16–5

- See also Page files
- as system dump file, *Vol 2*, 16–2, 16–8
- required location, *Vol 2*, 16–4
- requirement
 - location, *Vol 2*, 17–8

PAGEFILE n _NAME symbol, *Vol 2*, 16–29, 16–31

PAGEFILE n _SIZE symbol, *Vol 2*, 16–31

Page file process limit, *Vol 1*, 7–46

Page files

- as system dump file, *Vol 2*, 16–2, 16–3, 16–4, 16–8, 16–25
 - releasing dump from, *Vol 2*, 16–24
- changing size
 - recommended method, *Vol 2*, 16–22
- controlling location and size, *Vol 2*, 16–29
- creating
 - with AUTOGEN, *Vol 2*, 16–31
 - with SYSGEN, *Vol 2*, 16–33
- definition, *Vol 2*, 16–5
- deleting after creating a new version, *Vol 2*, 16–35
- displaying, *Vol 2*, 16–6
- fragmentation, *Vol 2*, 16–34
- freeing dump information from, *Vol 2*, 16–5, 16–24
- installing, *Vol 2*, 16–29
 - in system startup, *Vol 1*, 5–6; *Vol 2*, 16–6, 16–26, 16–27
 - secondary, *Vol 2*, 16–6
 - with SYPAGSWPFILES.COM procedure, *Vol 2*, 16–26
 - with SYSGEN, *Vol 2*, 16–25
- limiting usage of, *Vol 2*, 16–11
- location
 - specifying, *Vol 2*, 16–31
- message
 - indicating insufficient size, *Vol 2*, 16–10
 - indicating lack of installed, *Vol 1*, 5–6
- monitoring use of, *Vol 2*, 16–6
- mounting disk during system startup, *Vol 1*, 5–6; *Vol 2*, 16–26

Page files (cont'd)

- moving to improve performance, *Vol 2*, 17–8
 - on a satellite, *Vol 2*, 16–26
 - OpenVMS Cluster satellite node, *Vol 1*, 5–6
 - primary, *Vol 2*, 16–6
 - purging, *Vol 2*, 16–35
 - rebooting after creation, *Vol 2*, 16–34
 - releasing, *Vol 2*, 16–25
 - releasing dump from, *Vol 2*, 16–24
 - requirements
 - location, *Vol 2*, 16–4, 16–6, 17–8
 - size for saving dumps, *Vol 2*, 16–4
 - saving dump contents on reboot, *Vol 2*, 16–3
 - secondary, *Vol 2*, 16–6, 16–26
 - creating to improve system performance, *Vol 2*, 17–8
 - sizes
 - calculating, *Vol 2*, 16–6, 16–8, 16–9
 - changing, *Vol 2*, 16–28, 16–32, 16–33
 - current, *Vol 2*, 16–30
 - displaying AUTOGEN calculations, *Vol 2*, 16–29
 - required, *Vol 2*, 16–4, 16–9
 - required for saving system dumps, *Vol 2*, 16–8
 - specifying, *Vol 2*, 16–29, 16–31
 - when to increase size, *Vol 2*, 16–10
 - tasks for managing, *Vol 2*, 16–1
 - writing crash dumps to, *Vol 2*, 16–2
- ## Pagelets
- size of, *Vol 1*, 7–42
- ## Page overflow
- controlling, *Vol 1*, 14–52
- ## Pages
- size, *Vol 1*, 7–41
- ## Page setup modules, *Vol 1*, 14–52
- See also Device control modules
 - specifying forms, *Vol 1*, 14–47
- ## Page width and length
- specifying in forms, *Vol 1*, 14–47
- ## Paging, *Vol 2*, 16–5
- increased with vector processing, *Vol 2*, 28–7
- ## PAKs (Product Authorization Keys)
- installing before using VMSINSTAL.COM
 - command procedure, *Vol 1*, 3–4
 - loading in system startup, *Vol 1*, 5–5
 - preventing nodes from sharing, *Vol 1*, 3–5
- ## PAN command, *Vol 2*, 22–6
- ## PAN keypad function, *Vol 2*, 22–7
- ## Paper
- See Printer paper; Stock
- ## Paper jam
- pausing printer to fix, *Vol 1*, 14–70
- ## Parameter files
- See also System parameters
 - ALPHAVMSSYS.PAR (Alpha), *Vol 2*, 15–3
 - initializing parameters at boot time, *Vol 2*, 15–36

Parameter files

- ALPHAVMSSYS.PAR (Alpha) (cont'd)
 - role in the boot process, *Vol 1*, 4–2
 - booting with alternate, *Vol 1*, 4–8
 - changing
 - effects of, *Vol 2*, 15–34
 - with SYSGEN, *Vol 2*, 15–33
 - with SYSMAN, *Vol 2*, 15–28
 - common in an OpenVMS Cluster environment, *Vol 2*, 15–22
 - creating new
 - with SYSGEN, *Vol 2*, 15–35
 - default, *Vol 1*, 4–6; *Vol 2*, 15–3
 - including in MODPARAMS.DAT file, *Vol 2*, 15–22
 - limitation on error checking in, *Vol 2*, 15–19
 - MODPARAMS.DAT, *Vol 2*, 15–17
 - sample, *Vol 2*, 15–17
 - multiple, with AUTOGEN, *Vol 2*, 15–22
 - VAXVMSSYS.PAR (VAX), *Vol 2*, 15–3
 - initializing active parameters at boot time, *Vol 2*, 15–36
 - role in the boot process, *Vol 1*, 4–2
- ## Parameters
- passing to a startup command procedure, *Vol 1*, 5–20
 - system
 - See System parameters
- ## PARAMETERS command, *Vol 1*, 2–12
- See also System parameters
 - in SYSMAN, *Vol 2*, 15–26
 - summary, *Vol 2*, 15–30
- ## Parse styles
- related error messages, *Vol 1*, 10–6
- ## Partition, *Vol 2*, 25–2
- ## \$PASSWORD card, *Vol 1*, 8–39
- ## Password generators
- using to obtain initial password, *Vol 1*, 12–2
- ## Passwords
- conditions required for SYSMAN, *Vol 1*, 2–14
 - dual, *Vol 1*, 12–2
 - entering when logging in, *Vol 1*, 2–10
 - forced change, *Vol 1*, 12–4
 - for cluster access
 - modifying, *Vol 2*, 22–16
 - forgotten by user, *Vol 1*, 4–11, 7–23
 - history list, *Vol 1*, 12–3
 - how to preexpire, *Vol 1*, 12–3
 - InfoServer system, *Vol 2*, 25–7
 - initial, *Vol 1*, 12–2
 - length, *Vol 1*, 12–4
 - modifying system, *Vol 1*, 7–9
 - modifying user, *Vol 1*, 7–13
 - protecting, *Vol 1*, 12–5
 - secondary, *Vol 1*, 12–4
 - setting the expiration for, *Vol 1*, 12–4
 - with SYSMAN, *Vol 1*, 2–17

PATHWORKS

- startup restrictions, *Vol 2*, 25–14
- PATHWORKS for OpenVMS, *Vol 1*, 10–9
- Paused queue status, *Vol 1*, 14–59
- Pausing an output queue, *Vol 1*, 14–70
 - to align position of print for preprinted form, *Vol 1*, 14–70, 14–71
 - to change position of print, *Vol 1*, 14–70
- Pausing queues, *Vol 1*, 14–59
- PCFs (product configuration files), *Vol 1*, 3–16
 - command to create, *Vol 1*, 3–20
 - command to input, *Vol 1*, 3–26
 - command to output, *Vol 1*, 3–26
 - creating before installation, *Vol 1*, 3–19
 - creating during installation, *Vol 1*, 3–26
 - naming, *Vol 1*, 3–20
 - use existing, *Vol 1*, 3–26
- PCSI
 - See POLYCENTER Software Installation utility
- PCSI\$CONFIGURATION, *Vol 1*, 3–20
- PDB (product database)
 - definition, *Vol 1*, 3–16
 - how to use, *Vol 1*, 3–23
 - noncompliant products, *Vol 1*, 3–23
 - removing products from, *Vol 1*, 3–31
 - retrieving product information from, *Vol 1*, 3–30
 - tracking software dependencies, *Vol 1*, 3–24
 - volume label change, *Vol 1*, 3–30
- PDFs (product description files), *Vol 1*, 3–16
- Pending bad block log file
 - BADLOG.SYS, *Vol 2*, A–9
 - reserved file, *Vol 2*, A–9
- Pending job status
 - definition, *Vol 1*, 14–66
 - deleting a job in, *Vol 1*, 14–69
 - determining whether a job is in, *Vol 1*, 14–73
 - inducing with STOP/QUEUE/REQUEUE command, *Vol 1*, 14–68
 - requeuing a job in, *Vol 1*, 14–69
 - troubleshooting, *Vol 1*, 14–73
- Percent sign (%)
 - as wildcard character
 - using with tape volumes, *Vol 1*, 10–31
- Performance
 - disk, *Vol 1*, 9–35
 - effect of file extension on, *Vol 2*, 17–7
 - file system data caches, *Vol 2*, 18–1
 - importance of correct page file size, *Vol 2*, 16–9
 - importance of correct swap file size, *Vol 2*, 16–11
 - importance of sufficient hardware capacity, *Vol 2*, 17–5
 - improving
 - decompressing system libraries, *Vol 2*, 17–6
 - designing efficient applications, *Vol 2*, 17–4
 - disabling high-water marking, *Vol 2*, 17–6

Performance

- improving (cont'd)
 - encouraging batch processing, *Vol 2*, 17–3
 - for vector processing with batch queues, *Vol 2*, 28–7
 - installing frequently used images, *Vol 2*, 17–7, 17–17
 - manipulate memory, *Vol 2*, 17–8
 - moving page and swap files off system disk, *Vol 2*, 16–6
 - options, *Vol 2*, 17–6
 - reducing system disk I/O, *Vol 2*, 17–7
 - restricting the number of interactive users, *Vol 2*, 17–4
 - restricting user login hours, *Vol 2*, 17–4
 - setting RMS file extend parameters, *Vol 2*, 17–7
 - shared address data, *Vol 2*, 17–20
 - tuning the system, *Vol 2*, 17–4
 - with AUTOGEN feedback, *Vol 2*, 15–10
 - load balancing on public volumes, *Vol 1*, 9–15
 - managing
 - See also Tuning
 - distributing the work load, *Vol 2*, 17–3
 - evaluating tuning success, *Vol 2*, 17–6
 - hardware capacity, *Vol 2*, 17–5
 - installing images, *Vol 2*, 17–17
 - knowing your work load, *Vol 2*, 17–2
 - options, *Vol 2*, 17–6
 - system tuning, *Vol 2*, 17–4
 - tasks, *Vol 2*, 17–1
 - using MONITOR, *Vol 2*, 20–41
 - with vector processing, *Vol 2*, 28–7
 - work load management strategy, *Vol 2*, 17–3
 - monitoring, *Vol 2*, 17–2
 - tools used for, *Vol 2*, 17–2
 - reducing CPU and I/O time, *Vol 2*, 17–20
 - testing for disks, *Vol 1*, 9–14
 - vector processing, *Vol 2*, 28–4
- PGFLQUO process limit, *Vol 1*, 7–46
- PGFLQUOTA process limit
 - limiting page file usage with, *Vol 2*, 16–11
 - minimum recommended value, *Vol 2*, 16–11
 - value for efficient backups, *Vol 1*, 11–11
- Phase controller for UETP
 - See UETPHAS00.EXE file
- Physical memory
 - using shared address data to save, *Vol 2*, 17–20
- Physical system dump, *Vol 2*, 16–4
 - compared to selective dump, *Vol 2*, 16–12
- POLYCENTER Software Installation utility, *Vol 1*, 3–14
 - See also PRODUCT command
 - creating a PCF, *Vol 1*, 3–19
 - databases, *Vol 1*, 3–16
 - DCL interface, *Vol 1*, 3–14

- POLYCENTER Software Installation utility (cont'd)
 - installing layered products, *Vol 1*, 3–2
 - installing software, *Vol 1*, 3–24
 - managing installed software, *Vol 1*, 3–31
 - preliminary steps, *Vol 1*, 3–24
 - privileges required, *Vol 1*, 3–15
 - PRODUCT command
 - changes, *Vol 1*, 3–28
 - to run, *Vol 1*, 3–14
 - product configuration files, *Vol 1*, 3–16
 - registering products, *Vol 1*, 3–23
 - registering volume labels, *Vol 1*, 3–30
 - removing installed products, *Vol 1*, 3–31
 - starting the utility, *Vol 1*, 3–14
 - using a product database, *Vol 1*, 3–23
 - using to add or delete optional system files, *Vol 1*, 5–2
- Pooled resource, *Vol 1*, 7–3
- Port allocation classes
 - changing, *Vol 1*, 4–8
 - desassigning, *Vol 1*, 4–8
- Ports
 - setting up LAT, *Vol 1*, 14–14
- Position of print
 - aligning for preprinted forms, *Vol 1*, 14–70, 14–71
 - changing, *Vol 1*, 14–70
- PostScript printing, *Vol 1*, 14–9
- PRCLM process limit, *Vol 1*, 7–46
- Preventing autostart queues from starting, *Vol 1*, 14–62
- Primary bootstrap image
 - definition, *Vol 1*, 4–19
 - role in boot process, *Vol 1*, 4–2
- PRIMARY day
 - defining for accounts, *Vol 1*, 7–28
- Primary links, *Vol 1*, 10–47
- Primary page file, *Vol 2*, 16–6
 - See also PAGEFILE.SYS file; Page files
 - location requirement, *Vol 2*, 17–8
- Primary processors, *Vol 2*, 28–2
- Primary swap file, *Vol 2*, 16–6
 - See also SWAPFILE.SYS file; Swap files
- Primary Volume Descriptor (PVD)
 - using to mount ISO 9660 volumes, *Vol 1*, 9–49
- PRINT/DELETE command, *Vol 1*, 14–78
- PRINT command
 - bypassing symbiont formatting, *Vol 1*, 14–52
 - overriding default form-feed options with, *Vol 1*, 14–52
 - preventing users from executing, *Vol 1*, 14–62
 - processing of, *Vol 1*, 14–2
 - specifying a form, *Vol 1*, 14–47
 - specifying banner pages, *Vol 1*, 14–46
 - specifying job retention, *Vol 1*, 14–29
 - specifying scheduling priority, *Vol 1*, 14–36
- PRINT command (cont'd)
 - specifying setup and page setup modules, *Vol 1*, 14–55
- Printer paper
 - controlling with forms, *Vol 1*, 14–47, 14–51
 - pausing to align, *Vol 1*, 14–70
 - sheet feed, *Vol 1*, 14–50
 - specifying stock, *Vol 1*, 14–50
 - specifying width, *Vol 1*, 14–49
- Printers
 - controlling functions of, *Vol 1*, 14–52
 - LAT
 - See LAT software, printers
 - managing, *Vol 1*, 8–34
 - setting characteristics, *Vol 1*, 8–35, 14–14
 - in system startup, *Vol 1*, 5–12, 8–35
 - setting up before creating queues, *Vol 1*, 14–14
 - spooled, *Vol 1*, 14–14
 - definition, *Vol 1*, 8–35
 - despooling, *Vol 1*, 8–37
 - recommended use, *Vol 1*, 8–35
 - requirement, *Vol 1*, 8–35
 - sample configuration, *Vol 1*, 14–12
 - spooling, *Vol 1*, 8–36
 - testing spooling of, *Vol 1*, 8–37
 - troubleshooting, *Vol 1*, 14–72
- Print forms
 - See Forms
- Printing
 - distributed, *Vol 1*, 14–13
 - from applications using spooled printers, *Vol 1*, 8–35
 - job status, *Vol 1*, 14–66
 - PostScript, *Vol 1*, 14–9
 - remotely, *Vol 1*, 14–13
 - resuming at a specified position, *Vol 1*, 14–70, 14–71
- Printing files, requirements, *Vol 1*, 14–78
- Print jobs
 - See Batch jobs; Output jobs; Print queues
- Print queues
 - See Batch queues; Output queues
- Print symbionts
 - See Symbionts
- Priority, *Vol 1*, 7–2, 7–33
 - base, *Vol 1*, 7–2, 7–33
 - choosing for a batch queue, *Vol 1*, 14–33
 - effect of changing, *Vol 1*, 14–6
 - specifying for a batch queue, *Vol 1*, 14–21, 14–32
 - job scheduling, *Vol 1*, 14–36, 14–38
 - See also Job scheduling
 - changing for a job, *Vol 1*, 14–36, 14–38, 14–69
 - specifying for a job, *Vol 1*, 14–36, 14–67

- Private volumes, *Vol 1*, 9–15
- Privileged images, *Vol 2*, 17–19, 17–22
- Privileges
 - all, *Vol 1*, 7–4
 - allowing nonprivileged users to run programs that require, *Vol 2*, 17–22
 - changing in SYSMAN, *Vol 1*, 2–17
 - enhancement for installed files, *Vol 2*, 17–22
 - files, *Vol 1*, 7–4
 - for SYSTEM account, *Vol 1*, 2–10
 - GRPPRV, *Vol 1*, 12–8
 - process, *Vol 1*, 7–3
 - required for creating a CD-ROM, *Vol 1*, 10–42
 - required for UETP, *Vol 2*, 19–23
 - required to mount volumes, *Vol 1*, 9–22
- SECURITY
 - to mount volume with protected subsystems, *Vol 1*, 9–22
- SYSNAM, *Vol 1*, 9–23
- SYSPRV, *Vol 1*, 9–23
 - giving rights of system user, *Vol 1*, 12–8
- VOLPRO, *Vol 1*, 9–21
 - to mount volume as foreign, *Vol 1*, 9–22
- Problems
 - See also Troubleshooting
 - booting
 - fixing by booting with default parameter values, *Vol 1*, 4–9
 - fixing by booting with minimum startup, *Vol 1*, 4–15
 - hardware, *Vol 1*, 4–18
 - invalid boot block, *Vol 1*, 4–19
 - devices not recognized by the system, *Vol 1*, 8–8
 - forgotten password
 - fixing by booting without the UAF, *Vol 1*, 4–11
 - logging in, *Vol 1*, 4–10, 4–11
- Process dumps
 - disk quota, *Vol 2*, 16–35
 - privileged users, *Vol 2*, 16–36
 - security, *Vol 2*, 16–35, 16–37
- Processes
 - maintaining when disconnecting a terminal, *Vol 1*, 8–22
 - priority, *Vol 1*, 7–33
 - See also Priority, job scheduling; Priority, base
- Processing environments
 - multiprocessing
 - See Multiprocessing
 - vector processing
 - See Vector processing
- Processing job status, *Vol 1*, 14–66
- Process limits, *Vol 1*, 7–2
 - account jobs, *Vol 1*, 7–45
 - adjusting for vector processing, *Vol 2*, 28–7
- Process limits (cont'd)
 - AST queue, *Vol 1*, 7–43
 - CPU default
 - specifying a value for batch queues, *Vol 1*, 14–20, 14–32, 14–34
 - CPU maximum
 - specifying a value for batch queues, *Vol 1*, 14–21, 14–32, 14–34
 - CPU time, *Vol 1*, 7–44
 - detached process, *Vol 1*, 7–46
 - direct I/O count, *Vol 1*, 7–44
 - enqueue quota, *Vol 1*, 7–45
 - expiration, *Vol 1*, 7–45
 - jobwide logical name table, *Vol 1*, 7–45
 - open file, *Vol 1*, 7–45
 - page file, *Vol 1*, 7–46
 - process jobs, *Vol 1*, 7–46
 - setting, *Vol 1*, 7–42
 - subprocess creation, *Vol 1*, 7–46
 - system resources, *Vol 1*, 7–42
 - timer queue entry, *Vol 1*, 7–46
 - working set
 - quota, *Vol 1*, 7–47
- Processors
 - adding to a multiprocessing active set, *Vol 2*, 28–3
 - removing from a multiprocessing active set, *Vol 2*, 28–3
- Process quotas
 - See also Process limits
 - recommended values for backups, *Vol 1*, 11–11
 - setting before a backup, *Vol 1*, 11–10
- Product Authorization Keys
 - See PAKs
- PRODUCT command
 - See also POLYCENTER Software Installation Utility
 - DCL interface syntax, *Vol 1*, 3–14
- Product configuration file
 - See PCF
- Product dependencies, *Vol 1*, 3–24
- Product description file
 - See PDF
- PRODUCT INSTALL command, *Vol 1*, 3–24, 3–28
- Product list
 - obtaining, *Vol 1*, 3–7
 - VMSINSTAL.COM parameter, *Vol 1*, 3–6
- Profiles
 - in Mail, *Vol 1*, 7–41
 - in SYSMAN, *Vol 1*, 2–16
 - changing default directory, *Vol 1*, 2–18
 - changing privileges, *Vol 1*, 2–17
- Protected images, *Vol 2*, 17–19, 17–22, 17–23
- Protected subsystems, *Vol 1*, 7–16
 - enabling, *Vol 1*, 9–36
 - mounting volumes with, *Vol 1*, 9–22

Protection

- See also Protection codes; Security
 - ACL-based, *Vol 1*, 7–16, 7–35, 12–9
 - applied to system dump file, *Vol 2*, 16–5
 - applying to public disk volumes, *Vol 1*, 9–22
 - applying to queues, *Vol 1*, 14–23 to 14–26
 - assigning code when mounting a volume, *Vol 1*, 9–30
 - changing, *Vol 1*, 9–23
 - changing with PROTECTION qualifier, *Vol 1*, 9–25
 - default, *Vol 1*, 10–22
 - changing, *Vol 1*, 10–23
 - directory
 - changing with SET SECURITY/PROTECTION command, *Vol 1*, 10–27
 - specifying with /PROTECTION qualifier, *Vol 1*, 10–27
 - specifying with CREATE/DIRECTORY command, *Vol 1*, 10–27
 - disk file, *Vol 1*, 10–21
 - disk volumes, *Vol 1*, 9–23
 - displaying, *Vol 1*, 10–22
 - file, *Vol 1*, 10–19
 - default, *Vol 1*, 10–21, 10–23
 - directory, *Vol 1*, 10–19, 10–26
 - disk, *Vol 1*, 10–19
 - for public disk volumes, *Vol 1*, 9–22
 - ISO 9660-formatted media, *Vol 1*, 9–24
 - magnetic tape, *Vol 1*, 10–28
 - for interchange environments, *Vol 1*, 9–27
 - format for object, *Vol 1*, 12–8
 - mask, *Vol 1*, 9–24
 - mounting a volume with protected subsystems, *Vol 1*, 9–35
 - of devices, *Vol 1*, 8–6
 - overriding, *Vol 1*, 9–33
 - qualifier for SET VOLUME command, *Vol 1*, 9–25
 - UIC-based codes, *Vol 1*, 7–16
 - volumes, *Vol 1*, 10–19
 - disk, *Vol 1*, 9–21, 9–22
 - ISO 9660-formatted media, *Vol 1*, 9–24
 - standard-labeled, *Vol 1*, 9–26
 - tape, *Vol 1*, 9–22, 9–26, 10–38
 - VOLPRO privilege, *Vol 1*, 9–21
 - volume set
 - ISO 9660-formatted media, *Vol 1*, 9–24
- ## Protection checks
- MOUNT command
 - overriding, *Vol 1*, 9–30
- ## Protection codes
- access types, *Vol 1*, 12–8
 - changing, *Vol 1*, 10–23
 - format, *Vol 1*, 12–7
 - null access specification, *Vol 1*, 12–8
 - specifying, *Vol 1*, 10–21

Protocols

- LASTport, *Vol 2*, 25–5
- Proxy accounts
 - adding, *Vol 1*, 7–38
- Proxy authorization files
 - NET\$PROXY.DAT, *Vol 1*, 7–4, 7–38
 - NETPROXY.DAT, *Vol 1*, 7–4, 7–38
- Proxy database
 - logical name defining location, *Vol 1*, 5–8
 - managed by security server, *Vol 1*, 5–5
- Proxy logins
 - controlling system use, *Vol 1*, 7–40
- PRTSMB symbiont, *Vol 1*, 14–3
 - on LAT printers, *Vol 1*, 14–73
- PTFs (product text files), *Vol 1*, 3–16
- Public volumes
 - access to, *Vol 1*, 9–14
 - conditions for using, *Vol 1*, 9–14
 - creating with SYSTEM qualifier, *Vol 1*, 9–31
 - definition, *Vol 1*, 9–14
 - initializing, *Vol 1*, 9–18
 - guidelines, *Vol 1*, 9–19
 - load balancing, *Vol 1*, 9–15
 - mounting, *Vol 1*, 9–27, 9–28
 - in system startup, *Vol 1*, 5–11
 - mounting volume sets, *Vol 1*, 9–43
 - on large configurations, *Vol 1*, 9–14
 - on small configurations, *Vol 1*, 9–14
 - planning, *Vol 1*, 9–14
 - protecting, *Vol 1*, 9–22
 - setting protection, *Vol 1*, 9–25
 - testing disk performance, *Vol 1*, 9–14
- PURGE command
 - saving disk space, *Vol 1*, 9–68

Q

- Qlogic 1020ISP adapter
 - system tuning, *Vol 2*, 15–37
- QMAN\$MASTER.DAT file
 - See Master file of queue database
- QMAN\$MASTER logical name, *Vol 1*, 13–5
 - defining during system startup, *Vol 1*, 5–8
 - defining to reduce system disk I/O, *Vol 2*, 17–7
- QUANTUM system parameter, *Vol 2*, 28–8
- Queue characteristics
 - canceling, *Vol 1*, 14–31
 - defining, *Vol 1*, 14–30
 - deleting, *Vol 1*, 14–32
 - obtaining information about, *Vol 1*, 14–31, 14–59
 - problems
 - deleting, *Vol 1*, 14–76
 - mismatch, *Vol 1*, 14–75
 - specifying
 - on a queue, *Vol 1*, 14–21, 14–31
 - storage of in queue database, *Vol 1*, 14–20

Queue commands

- affected by multiple queue managers, *Vol 1*, 13–11
- creating queue database files, *Vol 1*, 13–7
- creating queues, *Vol 1*, 14–18
- deleting queues, *Vol 1*, 14–65
- displaying information about queue managers, *Vol 1*, 13–6
- displaying jobs, *Vol 1*, 14–65
- displaying queues, *Vol 1*, 14–58
- enabling autostart, *Vol 1*, 14–17
- managing banner pages, *Vol 1*, 14–46
- managing characteristics, *Vol 1*, 14–30
- managing device control libraries, *Vol 1*, 14–53
- managing forms and stock, *Vol 1*, 14–47
- managing queues, *Vol 1*, 14–57
- modifying jobs, *Vol 1*, 14–67
- modifying queues, *Vol 1*, 14–61
- pausing queues, *Vol 1*, 14–61
- relationship between starting and enabling autostart, *Vol 1*, 14–4
- setting UIC-based protection, *Vol 1*, 14–24
- showing UIC-based protection, *Vol 1*, 14–24
- specifying options, *Vol 1*, 14–20, 14–35
- starting queues
 - autostart, *Vol 1*, 14–17
 - nonautostart, *Vol 1*, 14–18
- starting the queue manager, *Vol 1*, 13–7
 - caution, *Vol 1*, 13–8
 - creating an additional queue manager, *Vol 1*, 13–12
 - restarting, *Vol 1*, 13–10
 - specifying failover list, *Vol 1*, 13–9
 - specifying name of queue manager, *Vol 1*, 13–12
 - specifying nodes to run the queue manager, *Vol 1*, 13–6
- stopping
 - all queues on a node, *Vol 1*, 13–10
 - queues, *Vol 1*, 14–17, 14–62
 - the queue manager, *Vol 1*, 13–9

Queue configurations

- sample batch queuing system, *Vol 1*, 14–5 to 14–7
- sample output queuing environment, *Vol 1*, 14–8 to 14–13

Queue databases

- See also Journal file of queue database; Master file of queue database; Queue file of queue database
- closing, *Vol 1*, 13–9
- creating files, *Vol 1*, 13–7
- default location, *Vol 1*, 13–4
- definition, *Vol 1*, 13–3
- detection of queue corruption, *Vol 1*, 14–80
- determining location, *Vol 1*, 13–9
 - of master file, *Vol 1*, 13–5
 - of queue and journal file, *Vol 1*, 13–6

Queue databases (cont'd)

- files comprising, *Vol 1*, 13–3
 - for multiple queue managers, *Vol 1*, 13–4, 13–11
 - naming convention, *Vol 1*, 13–4
 - function, *Vol 1*, 13–3
 - logical name defining location, *Vol 1*, 5–8
 - managing, *Vol 1*, 13–1
 - mounting the disk holding, *Vol 1*, 13–6
 - moving, *Vol 1*, 13–4
 - moving to reduce system disk I/O, *Vol 2*, 17–7
 - requirement in an OpenVMS Cluster environment, *Vol 1*, 13–5, 13–6
 - restoring damaged, *Vol 1*, 13–14
 - saving, *Vol 1*, 13–13
 - specifying location, *Vol 1*, 13–5
- ## Queue failover, *Vol 1*, 14–4
- ## Queue file of queue database, *Vol 1*, 13–3
- See also Queue databases
 - location, *Vol 1*, 13–6, 13–9
 - changing, *Vol 1*, 13–6
 - saving, *Vol 1*, 13–13
- ## Queue files
- detection of corruption, *Vol 1*, 14–80
- ## Queue forms
- See Forms
- ## Queue managers
- automatic restart, *Vol 1*, 13–3, 13–8, 13–10, 13–19
 - availability of, *Vol 1*, 13–9, 13–19
 - communication with job controller, *Vol 1*, 13–3
 - creating additional, *Vol 1*, 13–10
 - default name, *Vol 1*, 13–11
 - definition, *Vol 1*, 13–1
 - displaying information about, *Vol 1*, 13–6
 - failover, *Vol 1*, 13–3, 13–19
 - forcing, *Vol 1*, 13–10
 - list of nodes for, *Vol 1*, 13–9, 13–19
 - function, *Vol 1*, 13–1
 - improving performance, *Vol 1*, 13–15
 - limiting nodes that can run, *Vol 1*, 13–9
 - managing, *Vol 1*, 13–1
 - moving to another node in an OpenVMS Cluster system, *Vol 1*, 13–10
 - multiple, *Vol 1*, 13–3, 13–10
 - commands affected by, *Vol 1*, 13–11
 - managing queues with, *Vol 1*, 13–12
 - moving a queue to another queue manager, *Vol 1*, 13–12
 - restriction, *Vol 1*, 13–3
 - specifying queue manager name, *Vol 1*, 13–11, 13–12
 - naming, *Vol 1*, 13–12
 - relationship to queues, *Vol 1*, 13–1
 - restarting, *Vol 1*, 13–10
 - role in queuing process, *Vol 1*, 13–2, 14–2
 - print jobs, *Vol 1*, 14–3

Queue managers (cont'd)

- role in starting active autostart queues, *Vol 1*, 14-4
- specifying name of, *Vol 1*, 13-12
- specifying preferred order in which nodes run, *Vol 1*, 13-9, 13-10
- starting, *Vol 1*, 13-7
- stopping, *Vol 1*, 13-9
- troubleshooting, *Vol 1*, 13-16

Queue master file

- logical name defining location, *Vol 1*, 5-8

Queue names

- default for batch queue, *Vol 1*, 14-5
- default for print queue, *Vol 1*, 14-8
- defining, *Vol 1*, 14-15

Queue options, *Vol 1*, 14-20

- banner pages, *Vol 1*, 14-38
- controlling job performance and resources, *Vol 1*, 14-32
- controlling page and line overflow, *Vol 1*, 14-51
- device control libraries, *Vol 1*, 14-52
- forms, *Vol 1*, 14-46
- qualifiers for specifying, *Vol 1*, 14-20 to 14-23
- restricting access, *Vol 1*, 14-23
- retaining jobs, *Vol 1*, 14-27
- suppressing form feed, *Vol 1*, 14-52

Queues

See also Generic queues

- activating autostart, *Vol 1*, 14-16, 14-17
- allowing jobs to complete before shutdown, *Vol 1*, 14-63
- assigning a default form, *Vol 1*, 14-50
- assigning device control libraries, *Vol 1*, 14-54
- autostart, *Vol 1*, 14-4

See also Autostart feature; Autostart queues

- activating, *Vol 1*, 14-4
- on LAT queues, *Vol 1*, 14-4, 14-9
- starting, *Vol 1*, 14-4
- availability of, *Vol 1*, 14-4, 14-16
- batch
 - setting up for vector processing, *Vol 2*, 28-7

batch execution

See Batch queues

- canceling characteristics, *Vol 1*, 14-31
- changing DEFAULT form, *Vol 1*, 14-48
- changing options on, *Vol 1*, 14-61
- closing, *Vol 1*, 14-62

commands

See Queue commands

- creating, *Vol 1*, 14-15, 14-18
 - autostart execution, *Vol 1*, 14-15, 14-17
 - generic, *Vol 1*, 14-19
 - nonautostart execution, *Vol 1*, 14-18
- creating a form, *Vol 1*, 14-49
- defining a characteristic, *Vol 1*, 14-30

Queues (cont'd)

- deleting, *Vol 1*, 14-65
- deleting a job in, *Vol 1*, 14-69
- disabled, *Vol 1*, 14-80
- displaying information about, *Vol 1*, 14-58
- failover of, *Vol 1*, 14-16
- forms, *Vol 1*, 14-46
- gathering information with F\$GETQUI, *Vol 1*, 14-59
- generic batch, *Vol 1*, 14-6
 - See also Generic queues
- generic output, *Vol 1*, 14-10
 - See also Generic queues
- holding and releasing jobs, *Vol 1*, 14-67
- in an OpenVMS Cluster environment, *Vol 2*, 22-2
- initializing, *Vol 1*, 14-18
- managing with multiple queue managers, *Vol 1*, 13-12
- merging, *Vol 1*, 14-64
- modifying, *Vol 1*, 14-61
- monitoring, *Vol 1*, 14-58
- mounting a form on, *Vol 1*, 14-47, 14-51
- moving jobs from one queue to another, *Vol 1*, 14-64
- moving to another queue manager, *Vol 1*, 13-12
- on standalone workstations
 - batch, *Vol 1*, 14-5
- output execution
 - See Execution queues; Output queues
- pausing, *Vol 1*, 14-61
- print
 - See Output queues
- problems deleting, *Vol 1*, 14-76
- protection, *Vol 1*, 14-23 to 14-26
- reinitializing existing, *Vol 1*, 14-61
- requeuing
 - executing job, *Vol 1*, 14-68
 - pending job, *Vol 1*, 14-69
- simplifying startup, *Vol 1*, 14-4
- spooling printers, *Vol 1*, 8-36, 14-14
- stalled
 - status of, *Vol 1*, 14-74, 14-76
 - troubleshooting, *Vol 1*, 14-76
- standalone workstations
 - output, *Vol 1*, 14-8
- starting
 - autostart, *Vol 1*, 14-17
 - in startup command procedure, *Vol 1*, 14-17
 - in system startup, *Vol 1*, 5-12
 - nonautostart, *Vol 1*, 14-18
- startup command procedure, *Vol 1*, 5-12
 - sample, *Vol 1*, 14-18
- status, *Vol 1*, 14-58, 14-59
 - determining, *Vol 1*, 14-72

Queues (cont'd)

- stopped
 - status of, *Vol 1, 14-74*
- stopping, *Vol 1, 14-62*
 - abruptly, *Vol 1, 14-62*
 - all on a node, *Vol 1, 14-63*
 - before shutdown, *Vol 1, 14-63*
 - smoothly, *Vol 1, 14-62*
- types of, *Vol 1, 14-3*
 - output execution, *Vol 1, 14-3*

Queuing system

- See Batch and print queuing system

Quorum

- recalculating using IPC, *Vol 1, 9-78*

QUOTA.SYS file

- See Quota file

Quota file

- Analyze/Disk_Structure utility creates version, *Vol 2, A-9*
- contents, *Vol 1, 9-64*
- creating, *Vol 1, 9-65*
- deleting, *Vol 1, 9-67*
- disabling, *Vol 1, 9-67*
- QUOTA.SYS, *Vol 2, A-9*
- requirements for, *Vol 1, 9-65*
- reserved file, *Vol 2, A-9*
- UIC [0,0] in, *Vol 1, 9-63*
- updating, *Vol 1, 9-67*

Quotas

- See Disk quotas; Process limits; UAFs (user authorization files), resource limits
- disk
 - See Disk quotas
- process, *Vol 1, 7-2 to 7-47*
- to run UETP, *Vol 2, 19-23*

R

RAD

- assigning batch queues to, *Vol 1, 14-35*

RAID (redundant arrays of independent disks)

- volume shadowing support, *Vol 1, 11-43*

Read/write disk

- partitions, *Vol 2, 25-2*

Read access

- See also Access types, read
- for disk directory files, *Vol 1, 10-26*
- for disk files, *Vol 1, 10-21*
- granting through protection codes, *Vol 1, 12-8*

Read-ahead caching, *Vol 2, 18-8*

Read error

- if returned when booting, *Vol 1, 4-18*

Read operation

- See Access types, read

Real-time priority, *Vol 1, 7-33*

Rebooting

- after creation of system files, *Vol 2, 16-34*

Rebuilding volumes, *Vol 1, 8-5*

- Reconfiguration of software installation options, *Vol 1, 3-29*

Records

- blocking on tapes, *Vol 1, 9-11*
- size, *Vol 1, 9-33*

Recovering lost files, *Vol 1, 9-72*

Reducing I/O on system disk, *Vol 2, 17-7*

Redundant arrays of independent disks

- See RAID

Registering images with system version

- dependencies, *Vol 1, 5-23*

Reinitializing volumes, *Vol 1, 9-59*

Release Notes

- extracting from software products, *Vol 1, 3-25*
- on software products, *Vol 1, 3-25*

Release Notes option

- VMSINSTAL.COM, *Vol 1, 3-13*

Releasing a job, *Vol 1, 14-67*

Remote identifiers, *Vol 1, 12-11*

Remote InfoServer device

- BIND command, *Vol 2, 25-16*
- making available, *Vol 2, 25-16*
- mounting during system startup, *Vol 2, 25-16*

Remote log (archive) files, *Vol 2, 20-28*

Remote monitoring

- limitation, *Vol 2, 20-45*
- mixed-version OpenVMS Cluster systems, *Vol 2, 20-45*

Remote node database

- logical name defining location, *Vol 1, 5-8*

Remote nodes

- definition, *Vol 1, 2-13*

Remote printing, *Vol 1, 14-13*

Remote queue status, *Vol 1, 14-59*

Remote terminals, *Vol 1, 8-23*

Removing installed software products, *Vol 1, 3-31*

Repairing disk structure errors, *Vol 1, 9-71*

REPLY/ENABLE=SECURITY command

- enabling security operator terminals, *Vol 2, 20-30*

REPLY command

- canceling a user request, *Vol 1, 9-35; Vol 2, 20-20*
- closing current operator log file, *Vol 2, 20-18*
- disabling operator terminals, *Vol 1, 2-25; Vol 2, 20-19*
- enabling operator terminals, *Vol 1, 2-24; Vol 2, 20-18*
- specifying terminal, *Vol 2, 20-19*
- ensuring that correct volume is mounted, *Vol 1, 9-55*
- initializing tape, *Vol 1, 9-56*
- linking continuation volume to volume set, *Vol 1, 9-55*

- REPLY command (cont'd)
 - opening a new operator log file, *Vol 2*, 20–18, 20–23
 - putting request in wait state, *Vol 1*, 9–35
 - replying to requests, *Vol 1*, 2–26, 9–35
 - response not recorded in operator log file, *Vol 2*, 20–20
 - sending messages to users, *Vol 1*, 2–23
- Reporting errors
 - disk structure, *Vol 1*, 9–71
- Reporting on product dependencies, *Vol 1*, 3–24
- Reports
 - AUTOGEN, *Vol 2*, 15–23
 - SHOW CLUSTER
 - adding classes of data, *Vol 2*, 22–5
 - adding data, *Vol 2*, 22–8
 - adding fields of data, *Vol 2*, 22–5
 - changing default at startup, *Vol 2*, 22–14
 - command to modify, *Vol 2*, 22–10
 - compressing reports, *Vol 2*, 22–11
 - controlling displays, *Vol 2*, 22–6, 22–10
 - controlling with command procedures, *Vol 2*, 22–6, 22–16
 - moving display, *Vol 2*, 22–12
 - organization of, *Vol 2*, 22–5
 - panning, *Vol 2*, 22–7
 - scrolling, *Vol 2*, 22–14
- REQUEST command
 - identification number in operator log file, *Vol 2*, 20–20
 - recording request in operator log file, *Vol 2*, 20–20
 - sending requests to an operator, *Vol 1*, 2–26; *Vol 2*, 20–20
- Requeuing
 - executing job, *Vol 1*, 14–68
 - pending job, *Vol 1*, 14–69
- Reserved files, *Vol 2*, A–4
 - backup log file (BACKUP.SYS), *Vol 2*, A–9
 - bad block file (BADBLK.SYS), *Vol 2*, A–8
 - continuation file (CONTIN.SYS), *Vol 2*, A–9
 - index file (INDEXF.SYS), *Vol 2*, A–5
 - list of, *Vol 1*, 9–5
 - master file directory (MFD), *Vol 2*, A–8
 - pending bad block log file (BADLOG.SYS), *Vol 2*, A–9
 - quota file (QUOTA.SYS), *Vol 2*, A–9
 - storage bitmap file (BITMAP.SYS), *Vol 2*, A–8
 - volume security profile (SECURITY.SYS), *Vol 2*, A–9
 - volume set list file (VOLSET.SYS), *Vol 2*, A–9
- Reserved Memory Registry, *Vol 2*, 17–26
 - adding entries, *Vol 2*, 17–27
 - allocating reserved memory, *Vol 2*, 17–29
 - displaying reserved memory, *Vol 2*, 17–30
 - freeing reserved memory, *Vol 2*, 17–29
 - removing entries, *Vol 2*, 17–29
 - returning reserved memory, *Vol 2*, 17–31
- Reserved Memory Registry (cont'd)
 - using, *Vol 2*, 17–27, 17–31
- Reset modules, *Vol 1*, 14–52
 - See also Device control modules
- Resident images
 - installing (Alpha), *Vol 2*, 17–19
 - in system startup, *Vol 1*, 5–13
- Resource accounting, *Vol 2*, 21–1
- Resource affinity domain (RAD)
 - assigning batch queues to, *Vol 1*, 14–35
- Resource limits
 - See Process limits
- Resource types
 - limits, *Vol 1*, 7–3
- Resource use
 - producing reports of, *Vol 2*, 21–4
- Restarting
 - the queue manager, *Vol 1*, 13–10
- Restarting OPCOM (operator communication manager), *Vol 1*, 2–23
- Restore operations
 - with BACKUP, *Vol 1*, 11–29
- Restoring directories, *Vol 1*, 11–31
- Restoring files, *Vol 1*, 11–30
 - from an image backup, *Vol 1*, 11–45
 - from an incremental backup, *Vol 1*, 11–47
 - levels of directory access, *Vol 1*, 11–31
- Restoring the OpenVMS Registry, *Vol 2*, D–11
- Restoring the queue database, *Vol 1*, 13–14
- Restricted accounts, *Vol 1*, 7–31
- Restricting CPU time, *Vol 1*, 7–29
- RESTUSER.COM command procedure, *Vol 1*, 11–39
- Resuming printing of an output job, *Vol 1*, 14–70, 14–71
- Resuming queue status, *Vol 1*, 14–59
- Retained job status, *Vol 1*, 14–29
 - definition, *Vol 1*, 14–66
 - releasing jobs in, *Vol 1*, 14–68
- Retaining jobs in a queue
 - changing retention of a job, *Vol 1*, 14–29
- Rights databases, *Vol 1*, 5–8
- Rights list
 - adding to system, *Vol 1*, 10–14
- RIGHTSLIST.DAT files
 - default protection, *Vol 1*, 7–5
 - moving to reduce system disk I/O, *Vol 2*, 17–7
- RIGHTSLIST logical name
 - defining during system startup, *Vol 1*, 5–8
 - defining to reduce system disk I/O, *Vol 2*, 17–7
- RMS
 - access to files at record level, *Vol 1*, 10–28
 - assigns name to file, *Vol 1*, 9–5
- RMS_EXTEND_SIZE system parameter, *Vol 2*, 17–7

Root directory
 adding to an existing system disk, *Vol 1*, 2–32

Root volume
 definition, *Vol 1*, 9–44
 MFD on, *Vol 1*, 9–44

Routing
 definition, *Vol 2*, 23–5

RSM (Remote System Manager)
 startup restrictions, *Vol 2*, 25–14

RSX systems
 use of ODS-1 file structure with, *Vol 1*, 9–4

RT-11 volume
 block-addressable, *Vol 1*, 10–40

RTAn devices, *Vol 1*, 8–23

RUN command
 effect of installed images on, *Vol 2*, 17–18

RV60 optical disk drives
 supported by UETP, *Vol 2*, 19–7

RVN (relative volume number)
 See File identification, relative volume number

S

Satellites
 booting, *Vol 2*, 24–25
 cross-architecture booting, *Vol 2*, 24–28
 LAN MOP and DECnet MOP downline load service, *Vol 2*, 24–25
 migrating to LAN MOP downline load service, *Vol 2*, 24–26

SATELLITE_PAGE.COM command procedure,
Vol 1, 5–6; *Vol 2*, 16–26
 execution of during system startup, *Vol 1*, 5–5

SAVEDUMP system parameter, *Vol 2*, 16–3, 16–4, 16–8

Save sets, *Vol 1*, 11–7, 11–25
 backing up twice with /VERIFY qualifier, *Vol 1*, 11–67
 Files–11 disk, *Vol 1*, 11–8
 Get Save Set
 VMSINSTAL.COM option, *Vol 1*, 3–11
 listing the contents of, *Vol 1*, 11–19
 magnetic tape, *Vol 1*, 11–7
 multivolume, *Vol 1*, 11–31
 name, *Vol 1*, 11–25
 network, *Vol 1*, 11–8
 on multiple tapes, *Vol 1*, 11–22
 product
 storing temporarily during installation, *Vol 1*, 3–11
 sequential disk, *Vol 1*, 11–9
 types, *Vol 1*, 11–7

Saving the queue database, *Vol 1*, 13–13

Scalars
 definition, *Vol 2*, 28–4

SCBs (storage control blocks)
 in storage bitmap file, *Vol 2*, A–8

Scheduling, *Vol 1*, 7–2, 7–33
 See also Priority, job scheduling
 batch jobs, *Vol 1*, 14–36
 print jobs, *Vol 1*, 14–36

Scheduling priority
 for a batch or print job
 changing, *Vol 1*, 14–36

SCROLL command, *Vol 2*, 22–13

SCROLL keypad function, *Vol 2*, 22–7

SCSI (Small Computer Systems Interface), *Vol 1*, 11–43

SCSNODE system parameter, *Vol 2*, 25–11
 in DECnet-Plus networks, *Vol 2*, 27–4

SDA (System Dump Analyzer utility)
 See System Dump Analyzer utility

SDA CLUE
 dump off system disk, *Vol 2*, 16–20

SDA CLUE commands
 analyzing the system dump file, *Vol 2*, 16–2
 saving contents of system dump file, *Vol 2*, 16–22
 using to archive dump file information, *Vol 2*, 16–19
 using to collect dump file information, *Vol 2*, 16–19

Search lists
 priority of installed images, *Vol 2*, 17–18

Search path of Help Message database files, *Vol 1*, 5–28

Secondary Bootstrap image
 role in boot process, *Vol 1*, 4–2

SECONDARY day
 defining for accounts, *Vol 1*, 7–28

Secondary page and swap files
 See Page files; Swap files
 installing during system startup, *Vol 1*, 5–6

Secondary processors, *Vol 2*, 28–2

Sections
 global, *Vol 2*, 17–21
 pages, *Vol 2*, 17–21

Sectors
 Files–11
 definition, *Vol 2*, A–2

Security
 See also Protection; Security auditing; Security audit log files
 alarm messages, *Vol 2*, 20–22
 alarms
 enabling, *Vol 2*, 20–30
 enabling operator terminals, *Vol 2*, 20–22
 messages in operator log file, *Vol 2*, 20–22
 OPCOM alarm messages, *Vol 2*, 20–22
 OpenVMS Cluster, *Vol 2*, 22–16
 password management, *Vol 1*, 12–2
 protected subsystems, *Vol 1*, 9–36

Security (cont'd)

- protecting public disk volumes, *Vol 1*, 9–22
- protecting queues, *Vol 1*, 14–23
- protecting the system dump file, *Vol 2*, 16–5
- risk of compromise by installing images with privileges, *Vol 2*, 17–22
- specifying alarm events, *Vol 2*, 20–22
- SECURITY.AUDIT\$JOURNAL file
 - See Security audit log files
- SECURITY.SYS file
 - See Volume security profile
- Security auditing, *Vol 1*, 12–2, 12–13
 - archive file for each node in a cluster, *Vol 2*, 20–28
 - defining log file in system startup, *Vol 1*, 5–10
 - description, *Vol 2*, 20–26, 20–27
 - disabling events, *Vol 2*, 20–29
 - displaying using SHOW AUDIT command, *Vol 2*, 20–28
 - enabling operator terminal, *Vol 2*, 20–30
- Security audit log files
 - closing, *Vol 2*, 20–31
 - creating new version, *Vol 2*, 20–31
 - definition, *Vol 2*, 20–2
 - reviewing, *Vol 2*, 20–27
- Security management, *Vol 1*, 7–16
 - in SYSMAN on remote nodes, *Vol 1*, 2–16
 - protecting queues, *Vol 1*, 14–23 to 14–26
- Security operator terminals, *Vol 2*, 20–30
- Security server (SECURITY_SERVER)
 - starting, *Vol 1*, 5–5
- Security Server process
 - intrusion database, *Vol 1*, 12–7
 - network proxy database, *Vol 1*, 12–7
- SECURITY_AUDIT.AUDIT\$JOURNAL file
 - moving to reduce system disk I/O, *Vol 2*, 17–7
- SELECT command
 - in SHOW CLUSTER, *Vol 2*, 22–13
- Selective system dump, *Vol 2*, 16–4
 - compared to physical dump, *Vol 2*, 16–4, 16–12
 - order information is written in, *Vol 2*, 16–13
 - storing, *Vol 2*, 16–12
 - using to reduce size, *Vol 2*, 16–12
- SEQ (file sequence number)
 - See File identification, file sequence number
- Sequential-disk save set, *Vol 1*, 11–9
 - initializing, *Vol 1*, 11–9
 - mounting, *Vol 1*, 11–9
- Server MIB
 - defined, *Vol 2*, C–1
 - setting up, *Vol 2*, C–10
 - subagent objects, *Vol 2*, C–2
 - subagents, *Vol 2*, C–1
- Server queues, *Vol 1*, 14–3
- Server queue status, *Vol 1*, 14–59

Services

- bindings, *Vol 2*, 25–15
 - nodes, *Vol 2*, 26–4
 - password protection, *Vol 2*, 25–15
 - write protection, *Vol 2*, 25–15
- ## Sessions
- maintaining on more than one terminal, *Vol 1*, 8–22
 - maintaining when disconnecting a terminal, *Vol 1*, 8–22
- ## SET (Field) command
- SHOW CLUSTER, *Vol 2*, 22–11
- ## SET/STARTUP command
- in conversational boot, *Vol 1*, 4–14
- ## SET ACCOUNTING command
- controlling which resources are tracked, *Vol 2*, 21–3
 - starting up a new accounting file, *Vol 2*, 21–3
- ## SET ACL command
- modifying file characteristics, *Vol 1*, 10–24
- ## SET AUDIT command
- creating new version of security audit log file, *Vol 2*, 20–31
 - on local node only, *Vol 2*, 20–31
 - enabling security alarms, *Vol 2*, 20–30
 - enabling security auditing, *Vol 2*, 20–27
- ## SET command
- in conversational boot, *Vol 1*, 4–7
- ## SET DEVICE command
- spooling printers, *Vol 1*, 8–36
- ## SET DIRECTORY command
- changing directory characteristics, *Vol 1*, 10–27
 - limiting disk space consumed, *Vol 1*, 9–14
 - modifying disk file characteristics, *Vol 1*, 10–24
- ## SET ENTRY command, *Vol 1*, 14–67
- changing forms, *Vol 1*, 14–75
 - changing scheduling priority, *Vol 1*, 14–38
 - holding jobs, *Vol 1*, 14–67, 14–68
 - releasing jobs, *Vol 1*, 14–67
 - requeuing pending jobs, *Vol 1*, 14–69
 - specifying job retention, *Vol 1*, 14–29
- ## SET ENVIRONMENT command, *Vol 1*, 2–14
- ## SET FILE command
- example, *Vol 1*, 7–36
 - modifying disk file characteristics, *Vol 1*, 10–24
 - using to assign an alias, *Vol 1*, 10–25
 - using to modify file characteristics, *Vol 1*, 10–25
- ## SET FUNCTION command, *Vol 2*, 22–7
- in SHOW CLUSTER, *Vol 2*, 22–13
- ## SET HOST/LAT command, *Vol 2*, 26–2
- ## SET LOGINS command, *Vol 2*, 17–4
- ## SET MAGTAPE command, *Vol 1*, 8–38, 9–56
- ## SETPARAMS.DAT file, *Vol 2*, 15–18
- ## SET PRINTER command, *Vol 1*, 8–35, 14–14
- in system startup, *Vol 1*, 5–12

SET PROFILE command
in SYSMAN, *Vol 1*, 2–17

SET QUEUE command, *Vol 1*, 14–61
assigning a default form, *Vol 1*, 14–50
assigning characteristics, *Vol 1*, 14–31, 14–75
canceling characteristics, *Vol 1*, 14–31
controlling page overflow, *Vol 1*, 14–52
mounting a form, *Vol 1*, 14–51
setting UIC-based protection on queues, *Vol 1*, 14–24
specifying banner pages, *Vol 1*, 14–46
specifying job processing options, *Vol 1*, 14–35
specifying queue options with, *Vol 1*, 14–20
specifying reset modules, *Vol 1*, 14–53

SET QUORUM command
avoiding use of /CLUSTER with SYSMAN DO command, *Vol 2*, 22–18

SET RIGHTS_LIST command
using to restrict access to ODS-5 volumes, *Vol 1*, 10–14

SET SECURITY/ACL command
for vector capability (VAX), *Vol 2*, 28–8

SET SECURITY/PROTECTION command, *Vol 1*, 10–25
changing directory protection, *Vol 1*, 10–27
setting default protection, *Vol 1*, 10–23

SET SECURITY command
for queues, *Vol 1*, 14–24
modifying file characteristics, *Vol 1*, 10–24
using to prevent access to ODS-5 volumes, *Vol 1*, 10–15

SET SERVER REGISTRY_SERVER, *Vol 2*, D–8

SET TERMINAL command, *Vol 1*, 8–21, 14–14
despooling a printer, *Vol 1*, 8–37
determining characteristics of a LAT line, *Vol 1*, 8–24
enabling virtual terminals, *Vol 1*, 8–23
in system startup, *Vol 1*, 5–12, 8–22, 8–35
setting printer characteristics, *Vol 1*, 8–35

SET TIMEOUT command, *Vol 1*, 2–18

Setting system parameters
See Changing system parameters

Setting up printers, *Vol 1*, 8–21, 8–35, 14–14
in system startup, *Vol 1*, 5–12, 8–35
LAT, *Vol 1*, 14–14

Setting up terminals, *Vol 1*, 8–21, 8–35, 14–14
in system startup, *Vol 1*, 5–12, 8–22
using system parameters, *Vol 1*, 8–21

Setup module, *Vol 1*, 14–52
See also Device control modules
specifying in forms, *Vol 1*, 14–47

SET VOLUME /STRUCTURE_LEVEL=5 command, *Vol 1*, 9–37

SET VOLUME command
changing protection code, *Vol 1*, 9–25
disabling high-water marking, *Vol 2*, 17–7
encoding labels on volumes, *Vol 1*, 9–41

SET VOLUME command (cont'd)
modifying disk volume characteristics, *Vol 1*, 9–41
modifying file characteristics, *Vol 1*, 10–24
performing data checks, *Vol 1*, 9–41
specifying file retention periods, *Vol 1*, 9–69

Shadow sets
backing up, *Vol 1*, 11–43
mounting disks, *Vol 1*, 11–45
restoring, *Vol 1*, 11–50

Shareable images, *Vol 2*, 17–19, 17–21
assigning logical names for, *Vol 2*, 17–25

Shared address data
address data section, *Vol 2*, 17–21
assigns P1 space, *Vol 2*, 17–21
executable images can use, *Vol 2*, 17–21
image section, *Vol 2*, 17–20
Install utility determines content of address data section, *Vol 2*, 17–21
shared known images section, *Vol 2*, 17–20
using on Alpha systems, *Vol 2*, 17–20

Shared resource
definition, *Vol 2*, 22–2

Sheet-fed paper
specifying in forms, *Vol 1*, 14–47

Short report format
See Console report during UETP

SHOW/STARTUP command
in conversational boot, *Vol 1*, 4–14

SHOW ACCOUNTING command, *Vol 2*, 21–2

SHOW AUDIT command
to display event classes being audited, *Vol 2*, 20–28

SHOW CLUSTER command
See Show Cluster utility

Show Cluster utility (SHOW CLUSTER)
controlling displays, *Vol 2*, 22–6
defining keypad keys, *Vol 2*, 22–7
refreshing the screen, *Vol 2*, 22–10
reports, *Vol 2*, 22–5
controlling displays, *Vol 2*, 22–10
formatting, *Vol 2*, 22–6, 22–11, 22–14
initialization file, *Vol 2*, 22–10
using, *Vol 2*, 22–4

SHOW command, *Vol 1*, 10–16
in conversational boot, *Vol 1*, 4–7
in SYSGEN, *Vol 2*, 15–32
specifying languages, *Vol 1*, 6–18, 6–19
time and date formats, *Vol 1*, 6–17

SHOW CPU command, *Vol 2*, 28–3, 28–9

SHOW DEVICE/FULL command
to display type of volume, *Vol 1*, 10–3

SHOW DEVICES command, *Vol 1*, 8–2, 8–38, 10–20
checking mounted volumes, *Vol 1*, 9–52
determining the status of files, *Vol 1*, 9–58; *Vol 2*, 17–26
devices not shown, *Vol 1*, 8–8

SHOW DEVICES command (cont'd)
 examples, *Vol 1*, 8-3
 for ISO-9660 formatted devices, *Vol 1*, 8-5
 volume rebuild status, *Vol 1*, 8-5
 SHOW ENTRY command, *Vol 1*, 14-65
 SHOW ENVIRONMENT command, *Vol 1*, 2-14
 Showing system parameters
 with conversational boot, *Vol 1*, 4-7
 with SYSGEN, *Vol 2*, 15-32
 with SYSMAN, *Vol 2*, 15-27
 SHOW INTRUSION command, *Vol 1*, 12-7
 SHOW MEMORY command
 determining size of page, swap, and dump files,
 Vol 2, 16-30
 displaying page and swap files, *Vol 2*, 16-6
 monitoring page file use, *Vol 2*, 16-6
 showing the size of nonpaged pool, *Vol 2*, 28-7
 SHOW PROCESS command, *Vol 1*, 10-20; *Vol 2*,
 28-9
 SHOW PROFILE command
 in SYSMAN, *Vol 1*, 2-17
 SHOW PROTECTION command, *Vol 1*, 10-20,
 10-21
 SHOW QUEUE/MANAGERS command, *Vol 1*,
 13-6
 SHOW QUEUE command
 showing all jobs, *Vol 1*, 14-58
 showing batch jobs, *Vol 1*, 14-58
 showing brief information, *Vol 1*, 14-59
 showing complete information, *Vol 1*, 14-59
 showing completion status for jobs, *Vol 1*,
 14-27
 showing files associated with a job, *Vol 1*,
 14-59
 showing generic queues, *Vol 1*, 14-58
 showing jobs of a specified status, *Vol 1*, 14-58
 showing output execution queues, *Vol 1*, 14-58
 showing queue status, *Vol 1*, 14-58
 showing total number of jobs, *Vol 1*, 14-59
 SHOW QUOTA command, *Vol 1*, 9-66
 SHOW SECURITY command, *Vol 2*, 28-8
 for queues, *Vol 1*, 14-24
 SHOW SERVER REGISTRY_SERVER, *Vol 2*, D-7
 SHOW_CLUSTER\$INIT.COM command
 procedure, *Vol 2*, 22-14
 SHOW_CLUSTER\$INIT logical name, *Vol 2*,
 22-14
 Shutdown
 See SHUTDOWN.COM command procedure
 See System shutdown
 SHUTDOWN\$DECNET_MINUTES logical name,
 Vol 1, 4-28
 SHUTDOWN\$DISABLE_AUTOSTART logical
 name, *Vol 1*, 4-28, 14-64
 SHUTDOWN\$INFORM_NODES logical name,
 Vol 1, 4-28, 4-29
 SHUTDOWN\$MINIMUM_MINUTES logical
 name, *Vol 1*, 4-28
 SHUTDOWN\$QUEUE_MINUTES logical name,
 Vol 1, 4-29
 SHUTDOWN\$TIME logical name, *Vol 1*, 4-29
 SHUTDOWN\$VERBOSE logical name, *Vol 1*,
 4-29
 SHUTDOWN.COM command procedure, *Vol 1*,
 4-22
 See also System shutdown
 customizing, *Vol 1*, 4-28
 defining time before shutdown, *Vol 1*, 4-28
 example, *Vol 1*, 4-26
 executing with SYSMAN, *Vol 1*, 4-30
 how to use, *Vol 1*, 4-23
 options
 time of shutdown, *Vol 1*, 4-29
 order of events, *Vol 1*, 4-27
 reboot options, *Vol 1*, 4-24
 required privileges, *Vol 1*, 4-22
 when to use, *Vol 1*, 4-22
 Site-independent startup command procedure
 See STARTUP.COM command procedure
 Site-specific products
 startup database, *Vol 1*, 5-19
 Site-specific shutdown command procedure
 SYSHUTDWN.COM, *Vol 1*, 4-22
 Site-specific startup command procedure
 .COM version, *Vol 1*, 5-3
 creating your own, *Vol 1*, 5-2
 modifying, *Vol 1*, 5-2, 5-4
 order of execution, *Vol 1*, 4-5, 5-3
 required location, *Vol 1*, 5-4
 SATELLITE_PAGE.COM
 See SATELLITE_PAGE.COM command
 procedure
 SYCONFIG.COM
 See SYCONFIG.COM command procedure
 SYLOGICALS.COM
 See SYLOGICALS.COM command
 procedure
 SYPAGSWPFILES.COM
 See SYPAGSWPFILES.COM command
 procedure
 SYSECURITY.COM
 See SYSECURITY.COM command
 procedure
 SYSTARTUP_VMS.COM
 See SYSTARTUP_VMS.COM command
 procedure
 systemwide announcements, *Vol 1*, 5-14
 .TEMPLATE version, *Vol 1*, 5-3
 use in VMSKITBLD, *Vol 1*, 5-3
 Small computer systems interface
 See SCSI

SMISERVER process, *Vol 1*, 2–12
 attributes of, *Vol 1*, 2–17
 changing, *Vol 1*, 2–17
 starting, *Vol 1*, 2–12
 in system startup, *Vol 1*, 5–5

SMP (symmetric multiprocessing)
 See Multiprocessing

SMP_CPUS system parameter, *Vol 2*, 28–3, 28–6

SMTP (Simple Mail Transfer Protocol)
 specifying transport, *Vol 1*, 5–34

Software errors
 OPCOM failure, *Vol 1*, 2–23
 queue manager, *Vol 1*, 13–16
 when booting, *Vol 1*, 4–19

Software installation
 See Installing software

Software kits
 See Software products

Software license
 definition, *Vol 1*, 3–4

Software Performance Reports
 See SPRs

Software products
 See also Installing software
 consistent installation of, *Vol 1*, 3–19
 converting formats, *Vol 1*, 3–30
 copying to new locations, *Vol 1*, 3–30
 dependencies on other products, *Vol 1*, 3–24
 installing more than one, *Vol 1*, 3–25
 reconfiguring after installing, *Vol 1*, 3–29
 registering noncompliant software, *Vol 1*, 3–23
 release notes, *Vol 1*, 3–25
 removing after installing, *Vol 1*, 3–31
 tracking dependencies of, *Vol 1*, 3–24

Software Products Library
 using to update InfoServer software, *Vol 2*,
 25–3

SOGW user category abbreviation, *Vol 1*, 12–7

Sort/Merge utility (SORT/MERGE)
 optimizing batch queues for, *Vol 1*, 14–35

Source parameter
 VMSINSTAL.COM, *Vol 1*, 3–7

SPAWN function
 LAN Control Program (LANCP) utility, *Vol 2*,
 24–10

Spawn mode
 as execution mode for a startup procedure, *Vol 1*,
 5–20

Spoiled printers
 See Printers, spoiled

STABACKIT.COM command procedure, *Vol 1*,
 11–55, 11–57

Stalled job status, *Vol 1*, 14–66

Stalled queues
 status, *Vol 1*, 14–59

Standalone BACKUP
 booting, *Vol 1*, 11–56, 11–58
 building, *Vol 1*, 11–54, 11–57
 definition, *Vol 1*, 11–54
 qualifiers, *Vol 1*, 11–53
 relation to Backup utility, *Vol 1*, 11–54
 using to back up the system disk, *Vol 1*, 11–54,
 11–59, 11–62
 using to restore the system disk, *Vol 1*, 11–61

Standard time, *Vol 1*, 6–2
 changing to, *Vol 1*, 6–13
 automatically, *Vol 1*, 6–13

Star couplers
 definition of, *Vol 1*, 14–7

START/CPU command, *Vol 2*, 28–3, 28–6

START/QUEUE/MANAGER command, *Vol 1*,
 13–7, 13–10
 caution, *Vol 1*, 13–8
 creating an additional queue manager, *Vol 1*,
 13–12
 creating queue database files, *Vol 1*, 13–7
 specifying failover list, *Vol 1*, 13–9
 specifying name of queue manager, *Vol 1*,
 13–12
 specifying nodes to run the queue manager, *Vol 1*,
 13–6
 storage of, *Vol 1*, 13–8

START/QUEUE command, *Vol 1*, 14–18
 assigning a default form, *Vol 1*, 14–50
 assigning characteristics, *Vol 1*, 14–31
 canceling characteristics, *Vol 1*, 14–31
 controlling page overflow, *Vol 1*, 14–52
 mounting a form, *Vol 1*, 14–51
 resuming printing of a suspended job, *Vol 1*,
 14–70
 setting UIC-based protection on queues, *Vol 1*,
 14–24
 specifying autostart information, *Vol 1*, 14–17
 specifying banner pages, *Vol 1*, 14–46
 specifying job processing options, *Vol 1*, 14–35
 specifying queue options with, *Vol 1*, 14–20
 specifying reset modules, *Vol 1*, 14–53
 starting a generic queue, *Vol 1*, 14–19
 starting a nonautostart queue, *Vol 1*, 14–18

Starting InfoServer Client for OpenVMS, *Vol 2*,
 25–11

Starting job status, *Vol 1*, 14–66

Starting queues, *Vol 1*, 14–4
 autostart, *Vol 1*, 14–17
 in startup command procedure, *Vol 1*,
 14–17
 relationship to activating an autostart
 queue, *Vol 1*, 14–4
 nonautostart, *Vol 1*, 14–18
 in startup command procedure, *Vol 1*,
 14–18

- Starting queue status, *Vol 1*, 14–59
- Starting the LAT software
 - with LAT\$STARTUP.COM, *Vol 1*, 5–15; *Vol 2*, 26–16, 26–17
- Starting the OpenVMS Registry, *Vol 2*, D–5
 - manually, *Vol 2*, D–6
- Starting the queue manager, *Vol 1*, 13–7
 - initially, *Vol 1*, 13–7
 - restarting, *Vol 1*, 13–10
- STARTNET.COM command procedure, *Vol 1*, 5–16, 8–10
- STARTUP\$AUTOCONFIGURE_ALL symbol, *Vol 1*, 8–10
- STARTUP\$INTERACTIVE_LOGINS symbol, *Vol 1*, 5–16
- STARTUP\$STARTUP_LAYERED logical name, *Vol 1*, 5–19
- STARTUP\$STARTUP_VMS logical name, *Vol 1*, 5–19
- STARTUP.COM command procedure, *Vol 1*, 4–5, 5–3
 - configuring devices, *Vol 1*, 5–7, 8–7
 - definition, *Vol 1*, 5–2
 - description, *Vol 1*, 4–13
 - if it does not execute, *Vol 1*, 4–19
 - message indicating execution of, *Vol 1*, 4–6
 - tasks performed by, *Vol 1*, 4–5, 4–13; *Vol 2*, 17–18
- STARTUP command, *Vol 1*, 2–12
 - See also Startup database in SYSMAN, *Vol 1*, 5–18
- Startup command procedure
 - See also Site-specific startup command procedure; STARTUP.COM command procedure
 - booting without, *Vol 1*, 4–10
 - changing execution mode, *Vol 1*, 5–21
 - changing node restrictions, *Vol 1*, 5–21
 - changing startup phase, *Vol 1*, 5–21
 - creating your own, *Vol 1*, 8–22, 8–35
 - enabling a temporarily disabled, *Vol 1*, 5–22
 - known file lists, *Vol 1*, 5–13
 - modifiable
 - See Site-specific startup command procedure
 - node restriction, *Vol 1*, 5–20
 - passing parameters to, *Vol 1*, 5–20
 - preventing from executing, *Vol 1*, 5–22
 - temporarily, *Vol 1*, 5–22
 - required
 - See STARTUP.COM command procedure
- SATELLITE_PAGE.COM
 - See SATELLITE_PAGE.COM command procedure
- setting up output devices, *Vol 1*, 14–14
- site-independent
- Startup command procedure
 - site-independent (cont'd)
 - See also STARTUP.COM command procedure
 - specifying an alternate, *Vol 1*, 4–14 as the default, *Vol 1*, 4–15
 - site-specific
 - See Site-specific startup command procedure
 - specifying execution mode, *Vol 1*, 5–21
 - specifying node restrictions, *Vol 1*, 5–21
 - specifying startup phase, *Vol 1*, 5–21
 - starting queues, *Vol 1*, 14–18
- SYCONFIG.COM
 - See SYCONFIG.COM command procedure
- SYLOGICALS.COM
 - See SYLOGICALS.COM command procedure
- SYPAGSWPFILES.COM
 - See SYPAGSWPFILES.COM command procedure
- SYSECURITY.COM
 - See SYSECURITY.COM command procedure
- when errors prevent you from logging in, *Vol 1*, 4–10
- Startup database
 - adding files to, *Vol 1*, 5–21
 - changing information in, *Vol 1*, 5–21
 - definition, *Vol 1*, 5–18
 - deleting records in, *Vol 1*, 5–22
 - disabling files in, *Vol 1*, 5–22
 - reenabling disabled files in, *Vol 1*, 5–22
 - restriction, *Vol 1*, 5–21
 - showing contents of, *Vol 1*, 5–21
 - showing name of the target, *Vol 1*, 5–20
 - specifying the current, *Vol 1*, 5–20
- Startup phases
 - determining order of, *Vol 1*, 5–18
 - layered product, *Vol 1*, 5–18
 - END, *Vol 1*, 5–19
 - LPBEGIN, *Vol 1*, 5–19
 - LPBETA, *Vol 1*, 5–19
 - LPMAIN, *Vol 1*, 5–19
 - specifying, *Vol 1*, 5–21
 - operating system
 - BASEENVIRON, *Vol 1*, 5–4
 - CONFIGURE, *Vol 1*, 5–4
 - DEVICE, *Vol 1*, 5–4
 - INITIAL, *Vol 1*, 5–4
- Startup restrictions
 - InfoServer Client for OpenVMS software, *Vol 2*, 25–14
 - PATHWORKS, *Vol 2*, 25–14
 - RSM, *Vol 2*, 25–14
 - SYSMAN, *Vol 2*, 25–14

STARTUP SET OPTIONS command, *Vol 1*, 4–17

STARTUP SHOW OPTIONS command, *Vol 1*, 4–18

STARTUP_P1 system parameter, *Vol 1*, 4–16

STARTUP_P2 system parameter, *Vol 1*, 4–17

 SYSMAN startup logging, *Vol 1*, 4–17

Status of jobs

 See Job status

Status of queues

 See Queues, status

Status of volume rebuilds, *Vol 1*, 8–5

\$STATUS values

 accessing for uninstalled messages, *Vol 1*, 5–26

Stock

 See also Forms

 commands used with, *Vol 1*, 14–47

 mismatch, *Vol 1*, 14–47

 troubleshooting, *Vol 1*, 14–75

 specifying, *Vol 1*, 14–47

STOP/CPU command, *Vol 2*, 28–3, 28–6

STOP/QUEUE/MANAGER/CLUSTER command, *Vol 1*, 13–9

STOP/QUEUE/NEXT command, *Vol 1*, 14–17, 14–62

 with autostart queues, *Vol 1*, 14–62

STOP/QUEUE/RESET command, *Vol 1*, 14–17, 14–62

 with autostart queues, *Vol 1*, 14–62

STOP/QUEUE command, *Vol 1*, 14–61, 14–70

STOP/QUEUES/ON_NODE command, *Vol 1*, 13–10

 entering before shutting down a system, *Vol 1*, 14–63

 relationship to DISABLE

 AUTOSTART/QUEUES command, *Vol 1*, 14–63

Stopped queue status, *Vol 1*, 14–59

Stop pending queue status, *Vol 1*, 14–59

Stopping queues, *Vol 1*, 14–62

 abruptly, *Vol 1*, 14–62

 all queues on a node, *Vol 1*, 13–10, 14–63

 smoothly, *Vol 1*, 14–62

Stopping the queue manager, *Vol 1*, 13–9

Storage bitmap file

 BITMAP.SYS, *Vol 2*, A–8

 definition, *Vol 2*, A–8

 reserved file, *Vol 2*, A–8

 storage control blocks (SCBs), *Vol 2*, A–8

Storage control blocks

 See SCBs

StorageWorks RAID Array, *Vol 1*, 11–45

SUBMIT/DELETE command, *Vol 1*, 14–78

SUBMIT command

 preventing users from executing, *Vol 1*, 14–62

 processing of, *Vol 1*, 14–2

 specifying job retention, *Vol 1*, 14–29

 specifying scheduling priority, *Vol 1*, 14–36

SUBMON.COM command procedure

 sample, *Vol 2*, 20–43

 used with Monitor utility, *Vol 2*, 20–41

Subprocesses

 creating

 with the LANCP SPAWN command, *Vol 2*, 24–10

 subprocess creation limit, *Vol 1*, 7–2, 7–46

Substituting volumes, *Vol 1*, 9–35

Subsystem ACEs

 example, *Vol 1*, 12–11

Subsystems

 protected, *Vol 1*, 9–36

Supervisor mode

 logical names, *Vol 2*, 17–23

Suspended job status, *Vol 1*, 14–66

Suspending a job, *Vol 1*, 14–70

SVD (Supplementary Volume Descriptor)

 using to mount ISO 9660 volumes, *Vol 1*, 9–49

SWAPFILE.SYS file, *Vol 2*, 16–5

 See also Swap files

SWAPFILE_n_NAME symbol, *Vol 2*, 16–29, 16–31

SWAPFILE_n_SIZE symbol, *Vol 2*, 16–31

Swap files

 controlling location and size, *Vol 2*, 16–29

 creating

 with AUTOGEN, *Vol 2*, 16–31

 with SYSGEN, *Vol 2*, 16–33

 definition, *Vol 2*, 16–5

 deleting after creating a new version, *Vol 2*, 16–35

 displaying, *Vol 2*, 16–6

 fragmentation, *Vol 2*, 16–34

 installing, *Vol 2*, 16–29

 in system startup, *Vol 1*, 5–5, 5–6; *Vol 2*, 16–6, 16–26, 16–27

 secondary, *Vol 2*, 16–6

 with SYPAGSWPFILES.COM procedure, *Vol 2*, 16–26

 with SYSGEN, *Vol 2*, 16–25

 location, *Vol 2*, 16–6

 specifying, *Vol 2*, 16–31

 monitoring use of, *Vol 2*, 16–6, 16–11

 mounting disk during system startup, *Vol 1*, 5–6; *Vol 2*, 16–26

 moving to improve performance, *Vol 2*, 17–8

 on a satellite, *Vol 2*, 16–26

 OpenVMS Cluster satellite node, *Vol 1*, 5–6

 primary, *Vol 2*, 16–6

 purging, *Vol 2*, 16–35

 rebooting after creation, *Vol 2*, 16–34

 secondary, *Vol 2*, 16–6, 16–26

 creating to improve system performance, *Vol 2*, 17–8

 sizes

 calculating, *Vol 2*, 16–6, 16–11

 changing, *Vol 2*, 16–28, 16–32, 16–33

 current, *Vol 2*, 16–30

- Swap files
 - sizes (cont'd)
 - displaying AUTOGEN calculations, *Vol 2*, 16–29
 - increasing, *Vol 2*, 16–11
 - specifying, *Vol 2*, 16–29, 16–31
 - tasks for managing, *Vol 2*, 16–1
- SWAPFILES.COM command procedure
 - changing size of primary page, swap, and dump files, *Vol 2*, 16–32
- SWAPFILE symbol, *Vol 2*, 16–31
- Swapping
 - See also Swap files
 - to move information between physical memory and files stored on disk, *Vol 2*, 16–5
- SYCONFIG.COM command procedure, *Vol 1*, 5–3
 - AUTOGEN failure, *Vol 1*, 8–11
 - configuring devices, *Vol 1*, 8–7
 - in startup, *Vol 1*, 5–4
 - modifying to connect special devices, *Vol 1*, 5–7
 - modifying to mount disks early, *Vol 1*, 5–12
 - STARTUP\$AUTOCONFIGURE_ALL symbol, *Vol 1*, 8–10
- SYLOGICALS.COM command procedure, *Vol 1*, 5–3, 5–8
 - in system startup, *Vol 1*, 5–5
 - mounting the queue database disk, *Vol 1*, 13–6
 - redefining location of master file of queue database, *Vol 1*, 13–6
 - redefining location of system files, *Vol 2*, 17–8
 - to specify default state of operator log files, *Vol 2*, 20–24
- SYLOGIN.COM command procedure, *Vol 1*, 5–17, 7–20
 - ensuring execution, *Vol 1*, 7–13
 - sample systemwide, *Vol 1*, 7–21
- SYLOGOUT.COM command procedure, *Vol 1*, 7–22
- Symbionts, *Vol 1*, 14–61
 - communicating with, *Vol 1*, 14–70
 - default, *Vol 1*, 14–3
 - determining, *Vol 1*, 14–73
 - for LAT printers, *Vol 1*, 14–3, 14–73
 - function of, *Vol 1*, 13–2
 - LATSYM, *Vol 1*, 14–3, 14–73
 - PRTSMB on LAT printers, *Vol 1*, 14–73
 - role in processing print jobs, *Vol 1*, 14–3
 - user-written, *Vol 1*, 14–3
- Symbols
 - See also Logical names
 - defining
 - in MODPARAMS.DAT file, *Vol 2*, 15–19, 16–28
 - for system page, swap, and dump file sizes, *Vol 2*, 16–31 to 16–32
 - for system parameters, *Vol 2*, 15–19
 - NUM_ETHERADAPT, *Vol 2*, 15–21
 - NUM_NODES, *Vol 2*, 15–21
- Symbols (cont'd)
 - PAGEFILE n _NAME, *Vol 2*, 16–29
 - STARTUP\$AUTOCONFIGURE_ALL, *Vol 1*, 8–10
 - STARTUP\$INTERACTIVE_LOGINS, *Vol 1*, 5–16
 - SWAPFILE n _NAME, *Vol 2*, 16–29
- Symmetric multiprocessing
 - See Multiprocessing
- Symmetric vector processing configuration, *Vol 2*, 28–4
- SYPAGSWPFILES.COM command procedure, *Vol 1*, 5–3; *Vol 2*, 16–26
 - execution of during system startup, *Vol 1*, 5–5, 5–6
 - modifying to install page and swap files, *Vol 1*, 5–6
- SYS\$ANNOUNCE logical name, *Vol 1*, 5–14
- SYS\$AUDIT_SERVER_INHIBIT logical name, *Vol 1*, 5–10; *Vol 2*, 20–29
- SYS\$BATCH default queue name, *Vol 1*, 14–5
- SYS\$CONFIG.DAT file
 - system file, *Vol 1*, 8–12
 - used in autoconfiguration of devices, *Vol 1*, 8–12
- SYS\$DECDTM_INHIBIT logical name, *Vol 2*, 27–23
- SYS\$ERRLOG.DMP file, *Vol 2*, 16–3
- SYS\$ERRORLOG logical name
 - defining during system startup, *Vol 1*, 5–8
 - defining to reduce system disk I/O, *Vol 2*, 17–7
- SYS\$INPUT logical name, *Vol 2*, 19–32
- SYS\$JOURNAL logical name, *Vol 2*, 27–3
- SYS\$LANGUAGE logical name, *Vol 1*, 6–23
- SYS\$MONITOR logical name
 - defining during system startup, *Vol 1*, 5–8
 - defining to reduce system disk I/O, *Vol 2*, 17–7
- SYS\$OUTPUT logical name, *Vol 2*, 19–34
- SYS\$PRINT default queue name, *Vol 1*, 14–8
- SYS\$QUEUE_MANAGER.QMAN\$JOURNAL file
 - See Journal file of queue database
- SYS\$QUEUE_MANAGER.QMAN\$QUEUES file
 - See Queue file of queue database
- SYS\$QUEUE_MANAGER queue manager
 - as default queue manager, *Vol 1*, 13–11
- SYS\$STARTUP logical name, *Vol 1*, 4–5, 5–2, 5–18
- SYS\$SYLOGIN logical name, *Vol 1*, 5–17
- SYS\$TEST logical name, *Vol 2*, 19–4, 19–11, 19–20
- SYS\$UPDATE logical name
 - with VMSINSTAL.COM command procedure, *Vol 1*, 3–5
- SYS\$USER_CONFIG.DAT file
 - adding descriptions to, *Vol 1*, 8–12
 - rebuilding, *Vol 1*, 8–16
 - syntax rules, *Vol 1*, 8–13

- SYS\$USER_CONFIG.DAT file (cont'd)
 used in autoconfiguration of devices, *Vol 1*, 8–12
 user file, *Vol 1*, 8–12
- SYS\$WELCOME logical name, *Vol 1*, 5–15
- SYSBOOT
 See Conversational boot
 using to define port allocation classes, *Vol 1*, 4–8
- SYSBOOT.EXE file, *Vol 1*, 4–2
- SYSDUMP.DMP file, *Vol 2*, 16–2 to 16–5
 See also System dump file
 protection, *Vol 2*, 16–5
 required location, *Vol 2*, 16–4
- SYSECURITY.COM command procedure, *Vol 1*, 5–3, 5–10
 during system startup, *Vol 1*, 5–5
- SYSGEN, *Vol 1*, 10–14
- SYSGEN (System Generation utility)
 See System Generation utility
- SYSGEN parameters
 See System parameters
- SYSHUTDOWN.COM command procedure, *Vol 1*, 4–22
- SYSLOST.DIR file
 lost files in, *Vol 1*, 9–72
- SYSMAN
 See System Management utility
- SYSMANINI logical name, *Vol 1*, 2–19
- SYSMAN parameters
 See System parameters
- SYSMAN utility
 Reserved Memory Registry, *Vol 2*, 17–26
- SYSPRV privilege
 giving rights of system user, *Vol 1*, 12–8
- SYSTARTUP_VMS.COM command procedure, *Vol 1*, 5–3 to 5–17
 creating systemwide announcements, *Vol 1*, 5–14
 defining announcements in, *Vol 1*, 5–15
 defining location of systemwide login procedure, *Vol 1*, 5–17
 defining welcome messages in, *Vol 1*, 5–15
 disabling error checking in, *Vol 1*, 5–10
 editing to start DECnet VAX, *Vol 1*, 5–16
 enabling autostart, *Vol 1*, 5–12
 freeing page file of dump information, *Vol 2*, 16–5, 16–25
 installing known images, *Vol 1*, 5–13; *Vol 2*, 17–18
 installing resident images (Alpha), *Vol 1*, 5–13
 invoking LAT command procedures, *Vol 1*, 5–15
 invoking the System Dump Analyzer utility, *Vol 2*, 16–22
 limiting the number of interactive users, *Vol 1*, 5–16
- SYSTARTUP_VMS.COM command procedure (cont'd)
 making remote InfoServer disks available, *Vol 1*, 5–13; *Vol 2*, 25–16
 message indicating execution of, *Vol 1*, 4–6
 modifying to perform site-specific operations during system startup, *Vol 1*, 5–10
 operations performed in, *Vol 1*, 5–10
 purging the operator log file, *Vol 1*, 5–14
 saving contents of system dump file in, *Vol 2*, 16–22
 setting printer device characteristics, *Vol 1*, 5–12, 8–35
 setting terminal device characteristics, *Vol 1*, 5–12, 8–22
 special consideration about operator assistance for MOUNT command, *Vol 1*, 5–12
 starting InfoServer Client for OpenVMS, *Vol 1*, 5–13; *Vol 2*, 25–11
 starting queues, *Vol 1*, 5–12
 submitting batch jobs from, *Vol 1*, 5–14
- SYSTEM account
 changing passwords for security of, *Vol 1*, 2–10
 exercising caution with privileges, *Vol 1*, 2–10
 initial modification, *Vol 1*, 7–9
 in UAF, *Vol 1*, 7–6, 7–7
 logging in to, *Vol 1*, 2–10
 setting process quotas for efficient backups, *Vol 1*, 11–10
 using AUTHORIZE to modify process limits, *Vol 1*, 3–4
- System console
 ? message, *Vol 1*, 4–18
- System crash
 See Crash dumps; System failures
- System disks
 adding an alternate root directory, *Vol 1*, 2–32
 automatic mounting of, *Vol 1*, 5–11
 backing up, *Vol 1*, 5–35, 11–51, 11–59, 11–62
 backing up after installation, *Vol 1*, 3–9
 backing up for software installations, *Vol 1*, 3–3
 booting from an alternate, *Vol 1*, 4–4
 building with VMSKITBLD, *Vol 1*, 2–28
 completing a disk created with VMSKITBLD, *Vol 1*, 2–30
 configuring a system root added with VMSKITBLD, *Vol 1*, 2–34
 copying system files from, *Vol 1*, 2–27
 copying system files using VMSKITBLD, *Vol 1*, 2–31
 disk space needed to run UETP, *Vol 2*, 19–5
 fragmentation, *Vol 2*, 16–34
 installing software on alternate, *Vol 1*, 3–14
 moving files off to improve system performance, *Vol 2*, 17–7
 moving page and swap files off, *Vol 2*, 16–6
 not in volume sets, *Vol 1*, 9–44

System disks (cont'd)

- not ODS-5 volumes, *Vol 1*, 10–13
- quotas for, *Vol 1*, 9–65
- removing and adding optional system files, *Vol 1*, 5–1
- restoring, *Vol 1*, 11–61
- saving space by removing optional files, *Vol 1*, 5–1
- saving space on, *Vol 2*, 16–6
- testing with UETP, *Vol 2*, 19–34
- UETP error message, *Vol 2*, 19–24, 19–25
- UETP test image, *Vol 2*, 19–35

System dump

- designating order of processes in, *Vol 2*, 16–14
- off system disk (DOSD), *Vol 2*, 16–14, 16–15

System Dump Analyzer utility (SDA)

- analyzing the system dump file, *Vol 2*, 16–2
- in system startup, *Vol 1*, 5–13; *Vol 2*, 16–5, 16–19
- COPY command, *Vol 2*, 16–22
- determining cause of system failure, *Vol 2*, 16–19
- freeing dump information from the page file, *Vol 2*, 16–24
- invoked automatically on reboot, *Vol 2*, 16–19
- saving contents of system dump file, *Vol 2*, 16–22

System dump files

- See also SYSDUMP.DMP file
- analyzing, *Vol 2*, 16–19
- comparison of physical and selective system dumps, *Vol 2*, 16–12
- controlling location, *Vol 2*, 16–29
- copying, *Vol 2*, 16–5, 16–22, 16–23
- creating, *Vol 2*, 16–4
- default location, *Vol 2*, 16–2
- definition, *Vol 2*, 16–2
- deleting after creating a new version, *Vol 2*, 16–35
- freeing page file, *Vol 2*, 16–5
- installing, *Vol 2*, 16–29
- insufficient disk space, *Vol 2*, 16–12
- investigating cause of system failure, *Vol 2*, 16–19
- lack of, *Vol 2*, 16–2
- location, *Vol 2*, 16–4
- page file used to store, *Vol 2*, 16–3
- placing on another disk, *Vol 2*, 16–3
- protecting, *Vol 2*, 16–5, 16–25
- rebooting after creation, *Vol 2*, 16–34
- saving contents on reboot, *Vol 1*, 5–13; *Vol 2*, 16–22
- saving minimal information in, *Vol 2*, 16–12
- sizes, *Vol 2*, 16–4, 16–31
 - calculating, *Vol 2*, 16–2, 16–7
 - changing, *Vol 2*, 16–12, 16–28, 16–33
 - controlling, *Vol 2*, 16–29

System dump files

- sizes (cont'd)
 - displaying AUTOGEN calculations, *Vol 2*, 16–29
 - manipulating, *Vol 2*, 16–30
 - minimizing, *Vol 2*, 16–12
- storing selective portions of memory, *Vol 2*, 16–12

System failures, *Vol 2*, 16–2

- See also Crash dumps; System dump files
- determining cause, *Vol 2*, 16–2, 16–19
- reporting to Compaq support representative, *Vol 2*, 16–19
- saving contents of system dump file after, *Vol 1*, 5–13; *Vol 2*, 16–22
- writing of system dump file, *Vol 2*, 16–2

System files

- duplicating using VMSKITBLD, *Vol 1*, 2–27
- moving off system disk to improve performance, *Vol 2*, 17–7
- on public volumes, *Vol 1*, 9–14
- optional
 - adding or deleting, *Vol 1*, 5–2
- System Generation utility (SYSGEN)
 - and version checking, *Vol 1*, 5–23
- AUTOCONFIGURE command (VAX)
 - in system startup, *Vol 1*, 5–7
- changing page, swap, and dump file sizes, *Vol 2*, 16–28, 16–33
- changing system parameter file with, *Vol 2*, 15–34
- changing system parameters, *Vol 2*, 15–33, 15–34, 20–21
 - See also System parameters, changing
- configuring devices
 - in system startup, *Vol 1*, 5–7
- CONNECT command (VAX), *Vol 1*, 8–8
 - in system startup, *Vol 1*, 5–7
- converting parameters for use with AUTOGEN, *Vol 2*, 15–5
- CREATE command, *Vol 2*, 16–33
- creating a new system parameter file, *Vol 2*, 15–35
- creating page, swap, and dump files, *Vol 2*, 16–33
- INSTALL command, *Vol 2*, 16–25
 - in SYPAGSWPFILES.COM, *Vol 2*, 16–27
 - in system startup, *Vol 1*, 5–7
- installing page, swap, and dump files
 - in system startup, *Vol 1*, 5–7
- LOAD command (VAX), *Vol 1*, 8–8
- managing system parameters, *Vol 2*, 15–4
- operator log messages, *Vol 2*, 20–21
- showing system parameters, *Vol 2*, 15–32

System hangup, *Vol 2*, 19–21, 19–30

- System libraries
 - decompressing, *Vol 2*, 17–6
 - expanding, *Vol 2*, 17–8, 17–13
 - LIBDECOMP.COM command procedure, *Vol 2*, 17–9
 - listing those known to LIBDECOMP.COM, *Vol 2*, 17–12
 - reducing, *Vol 2*, 17–15
- System management
 - centralizing with SYSMAN, *Vol 1*, 2–10
 - clearing counters, *Vol 2*, 24–30
 - creating access control lists (ACLs), *Vol 1*, 12–9
 - deleting device information, *Vol 2*, 24–22
 - deleting node information, *Vol 2*, 24–25
 - disabling MOP downline load service, *Vol 2*, 24–28
 - displaying device information, *Vol 2*, 24–21
 - displaying LAN device configurations, *Vol 2*, 24–10
 - displaying LAN device parameters, *Vol 2*, 24–11
 - displaying node information, *Vol 2*, 24–23
 - displaying OPCOM messages, *Vol 2*, 24–30
 - displaying status and counters, *Vol 2*, 24–28, 24–29
 - enabling MOP downline load service, *Vol 2*, 24–28
 - environment, *Vol 1*, 2–13 to 2–16
 - Load Trace facility, *Vol 2*, 24–30
 - on multiple nodes, *Vol 1*, 2–14
 - stopping the LANACP utility, *Vol 2*, 24–7
 - tasks
 - clusterwide management, *Vol 2*, 22–3
- System Management utility (SYSMAN)
 - accessing disks, *Vol 1*, 9–65
 - adding startup files to a startup database, *Vol 1*, 5–21
 - ALF commands, *Vol 1*, 7–34
 - authorization checks in, *Vol 1*, 2–16
 - changing privileges in, *Vol 1*, 2–17
 - changing system parameters
 - active values, *Vol 2*, 15–29
 - command verification in, *Vol 1*, 2–18
 - configuring devices (Alpha)
 - in system startup, *Vol 1*, 5–7
 - converting parameters for use with AUTOGEN, *Vol 2*, 15–5
 - creating command procedures for, *Vol 1*, 2–19
 - deleting startup files, *Vol 1*, 5–22
 - disabling startup files, *Vol 1*, 5–22
 - DISKQUOTA commands, *Vol 1*, 9–65
 - disk quotas with, *Vol 1*, 9–63
 - Disk Quota utility, *Vol 1*, 9–65
 - DO command, *Vol 1*, 2–18
 - enabling remote systems to execute commands, *Vol 1*, 2–12
 - enabling startup files, *Vol 1*, 5–22
 - features of, *Vol 1*, 2–10
 - System Management utility (SYSMAN) (cont'd)
 - how commands execute, *Vol 1*, 2–11
 - initialization file, *Vol 1*, 2–19
 - IO AUTOCONFIGURE command (Alpha)
 - in system startup, *Vol 1*, 5–7
 - IO CONNECT command (Alpha), *Vol 1*, 8–9
 - in system startup, *Vol 1*, 5–7
 - IO LOAD (Alpha), *Vol 1*, 8–9
 - IO REBUILD/VERIFY command, *Vol 1*, 8–17
 - loading licenses with, *Vol 1*, 3–4
 - management environment, *Vol 1*, 2–13
 - managing an OpenVMS Cluster, *Vol 2*, 22–16 to 22–20
 - managing startup, *Vol 1*, 5–18
 - managing system parameters, *Vol 2*, 15–4, 15–26
 - modifying the system parameter file, *Vol 2*, 15–28
 - PARAMETERS command, *Vol 2*, 15–26
 - privileges required, *Vol 1*, 2–11
 - profile, *Vol 1*, 2–16
 - adjusting, *Vol 1*, 2–17
 - rights limitation, *Vol 1*, 2–11
 - showing system parameters, *Vol 2*, 15–27
 - showing the contents of a startup database, *Vol 1*, 5–21
 - showing the name of the target startup database, *Vol 1*, 5–20
 - shutdown, *Vol 1*, 4–30
 - SMISERVER process, *Vol 1*, 2–12
 - specifying the current startup database, *Vol 1*, 5–20
 - STARTUP command, *Vol 1*, 5–18
 - startup logging, *Vol 1*, 4–17
 - startup restrictions, *Vol 2*, 25–14
 - SYSMAN\$NODE_TABLE logical name table
 - creating, *Vol 1*, 2–15
 - timeout periods, *Vol 1*, 2–18
 - usage restriction, *Vol 1*, 2–11
 - use of passwords, *Vol 1*, 2–14, 2–17
 - using logical names in, *Vol 1*, 2–15
 - using to centralize system management, *Vol 1*, 2–10
 - using to rebuild SYS\$USER_CONFIG.DAT file, *Vol 1*, 8–17
- System messages
 - See also Messages
 - using when managing a system, *Vol 1*, 2–5
- System parameters
 - See also Parameter files
 - ACP cache system, *Vol 1*, 9–57
 - active values, *Vol 2*, 15–3, 15–26
 - affected by AUTOGEN calculations, *Vol 2*, 15–10
 - ALPHAVMSSYS.PAR file (Alpha), *Vol 2*, 15–3
 - automatic setting by AUTOGEN, *Vol 2*, 15–3
 - booting with default, *Vol 1*, 4–9
 - categories by function, *Vol 2*, 15–2

System parameters (cont'd)

- changing
 - checking AUTOGEN's settings, *Vol 2*, 15-10
 - editing MODPARAMS.DAT file, *Vol 2*, 15-18
 - recommended method, *Vol 2*, 15-4, 15-18
 - specifying values in MODPARAMS.DAT file, *Vol 2*, 15-4
 - with AUTOGEN, *Vol 2*, 15-18
 - with conversational boot, *Vol 1*, 4-4, 4-7; *Vol 2*, 15-36
 - with SYSGEN, *Vol 2*, 15-33, 15-34
 - with SYSMAN, *Vol 2*, 15-28, 15-29
- checking before using VMSINSTAL.COM command procedure, *Vol 1*, 3-3
- controlling
 - increasing, *Vol 2*, 15-19
 - in MODPARAMS.DAT file, *Vol 2*, 15-4, 15-18
 - specifying absolute values, *Vol 2*, 15-19
 - specifying maximum values, *Vol 2*, 15-20
 - specifying minimum values, *Vol 2*, 15-20
 - with ADD_ prefix, *Vol 2*, 15-19
 - with MAX_ prefix, *Vol 2*, 15-20
 - with MIN_ prefix, *Vol 2*, 15-20
- creating a new parameter file
 - with SYSGEN, *Vol 2*, 15-35
- current values, *Vol 2*, 15-3, 15-26
- default values, *Vol 2*, 15-3
- definition, *Vol 2*, 15-2
- DUMPBUG, *Vol 2*, 16-4
- DUMPSTYLE, *Vol 2*, 16-4, 16-12
- dynamic, *Vol 2*, 15-3
- effect on other parameters, *Vol 2*, 15-4
- ERLBUFFERPAGES, *Vol 2*, 16-7
- ERRORLOGBUFFERS, *Vol 2*, 16-7
- file extensions, *Vol 2*, 17-7
- GBLPAGES, *Vol 2*, 17-21
- GBLSECTIONS, *Vol 2*, 17-21
- initialization at boot time, *Vol 2*, 15-36
- in memory
 - See Active system parameters
- managing, *Vol 2*, 15-1
- MODPARAMS.DAT file
 - See MODPARAMS.DAT file
- MULTIPROCESSING, *Vol 2*, 28-3
- MVTIMEOUT, *Vol 1*, 9-75, 9-77
- NPAGEDYN, *Vol 2*, 28-7
- on disk, *Vol 2*, 15-3
 - See also Current system parameters
 - in ALPHAVMSSYS.PAR file (Alpha), *Vol 2*, 15-36
 - in VAXVMSSYS.PAR file (VAX), *Vol 2*, 15-36
- parameter files
 - See Parameter files

System parameters (cont'd)

- QUANTUM, *Vol 2*, 28-8
- recommended method for changing, *Vol 2*, 15-4, 15-5
- RMS_EXTEND_SIZE, *Vol 2*, 17-7
- SAVEDUMP, *Vol 2*, 16-3, 16-4
- SCSNODE, *Vol 2*, 25-11
- showing
 - with conversational boot, *Vol 1*, 4-4, 4-7; *Vol 2*, 15-36
 - with SYSGEN, *Vol 2*, 15-32
 - with SYSMAN, *Vol 2*, 15-27
- SMP_CPUS, *Vol 2*, 28-3, 28-6
- STARTUP_P1, *Vol 1*, 4-16
- STARTUP_P2, *Vol 1*, 4-17
- storing your changes for use with AUTOGEN, *Vol 2*, 15-5
- symmetric multiprocessing, *Vol 2*, 28-3
- TAPE_MVTIMEOUT, *Vol 1*, 9-75, 9-77
- TTY_DEFCHAR, *Vol 1*, 8-21
- TTY_DEFCHAR2, *Vol 1*, 8-21, 8-23
- types of, *Vol 2*, 15-2
 - dynamic, *Vol 2*, 15-2
 - general, *Vol 2*, 15-2
 - major, *Vol 2*, 15-2
- UAFALTERNATE, *Vol 1*, 4-12
- user definable, *Vol 2*, 15-3
- VAXVMSSYS.PAR file (VAX), *Vol 2*, 15-3
- VBN_CACHE_S, *Vol 2*, 18-3
- VCC_FLAGS, *Vol 2*, 18-3
- VCC_MAXSIZE, *Vol 2*, 18-11
- VCC_MAX_CACHE, *Vol 2*, 18-5
- VCC_MAX_IO_SIZE, *Vol 2*, 18-6
- VCC_PTES, *Vol 2*, 18-11
- VCC_READAHEAD, *Vol 2*, 18-8
- vector processing, *Vol 2*, 28-5
- VECTOR_MARGIN, *Vol 2*, 28-8
- VECTOR_PROC, *Vol 2*, 28-5
- VIRTUALPAGECNT, *Vol 1*, 14-35
- when incorrect values prevent the system from booting, *Vol 1*, 4-9
- WSMAX, *Vol 1*, 14-34; *Vol 2*, 28-7

System passwords, *Vol 1*, 12-3

- dictionary of, *Vol 1*, 12-3

Systems

- downline loading with InfoServer, *Vol 2*, 25-2

System shutdown

- See also SHUTDOWN.COM command procedure
- adjusting quorum when shutting down a node, *Vol 1*, 4-24
- after software installation, *Vol 1*, 3-9
- allowing batch and print jobs to complete before, *Vol 1*, 14-63
- caution about timing of system halt, *Vol 2*, 16-2
- checking for existence of system files before, *Vol 1*, 4-24

System shutdown (cont'd)

- customizing, *Vol 1*, 4-28
 - with SYSHUTDOWN.COM command procedure, *Vol 1*, 4-22
 - defining the minimum number of minutes before shutdown, *Vol 1*, 4-28
 - emergency procedures
 - console, *Vol 1*, 4-22
 - OPCCRASH, *Vol 1*, 4-22, 4-31
 - for an entire OpenVMS Cluster, *Vol 1*, 4-24
 - notification of, *Vol 1*, 4-29
 - options
 - automatic reboot, *Vol 1*, 4-24
 - manual reboot, *Vol 1*, 4-24
 - specifying time interval between DISABLE AUTOSTART/QUEUES and STOP/QUEUES/ON_NODE commands, *Vol 1*, 14-64
 - time of shutdown, *Vol 1*, 4-29
 - orderly, *Vol 1*, 4-22
 - order of events, *Vol 1*, 4-27
 - procedures for performing, *Vol 1*, 4-22
 - saving AUTOGEN feedback data, *Vol 1*, 4-24
 - SHUTDOWN.COM, *Vol 1*, 4-22
 - example, *Vol 1*, 4-26
 - how to use, *Vol 1*, 4-23
 - when to use, *Vol 1*, 4-22
 - stopping queues before, *Vol 1*, 14-63
 - with SYSMAN, *Vol 1*, 4-30
- ## System startup
- See also Booting; Site-specific startup command procedure; Startup command procedure; Startup phases
- analyzing a crash dump, *Vol 1*, 5-13
 - assigning systemwide logical names, *Vol 1*, 5-8
 - booting with minimum, *Vol 1*, 4-15
 - CONFIGURE phase, *Vol 1*, 8-7
 - configuring devices, *Vol 1*, 4-5, 8-7
 - special (Alpha), *Vol 1*, 5-7
 - special (VAX), *Vol 1*, 5-7
 - controlling when booting, *Vol 1*, 4-13
 - creating systemwide announcements, *Vol 1*, 5-14
 - databases, *Vol 1*, 5-18
 - defining location of systemwide login procedure, *Vol 1*, 5-17
 - definition of logical names, *Vol 1*, 5-4, 5-5
 - description, *Vol 1*, 4-13
 - determining order of phases, *Vol 1*, 5-18
 - displaying startup commands as they execute, *Vol 1*, 4-16
 - enabling autostart, *Vol 1*, 5-12
 - enabling operator console, *Vol 1*, 5-5
 - enabling operator log file, *Vol 1*, 5-5
 - events, *Vol 1*, 4-5
 - order of, *Vol 1*, 5-4
 - possibility of future change in order, *Vol 1*, 5-5

System startup (cont'd)

- execution of AUTOCONFIGURE command, *Vol 1*, 5-4
- execution of login procedures, *Vol 1*, 5-17
- execution of site-specific startup command procedures, *Vol 1*, 5-5
- freeing dump information from page file, *Vol 2*, 16-5
- in an emergency
 - with default system parameters, *Vol 1*, 4-9
 - without startup and login procedures, *Vol 1*, 4-10
 - without the UAF, *Vol 1*, 4-11
- installing images, *Vol 1*, 5-4; *Vol 2*, 17-18
- installing page and swap files, *Vol 1*, 5-5, 5-6; *Vol 2*, 16-6, 16-26, 16-27
- limiting the number of interactive users, *Vol 1*, 5-16
- LMF database, *Vol 1*, 5-5
- loading of device drivers, *Vol 1*, 5-4
- loading of licenses, *Vol 1*, 5-5
- loading of Product Authorization Keys (PAKs), *Vol 1*, 5-5
- location of files used in, *Vol 1*, 4-5
- logging with SYSMAN, *Vol 1*, 4-17
- making remote InfoServer devices available, *Vol 2*, 25-16
- making remote InfoServer disks available, *Vol 1*, 5-13
- managing with SYSMAN, *Vol 1*, 5-18
- messages
 - indicating execution of site-specific startup, *Vol 1*, 4-6
 - indicating lack of installed page file, *Vol 1*, 5-6
- mounting disk for page and swap files, *Vol 1*, 5-6
- mounting the queue database disk, *Vol 1*, 13-6
- not extended file name, *Vol 1*, 10-14
- performing site-specific operations, *Vol 1*, 5-10
- phases
 - See Startup phases
- purging the operator log file, *Vol 1*, 5-14
- saving contents of system dump file, *Vol 1*, 5-13; *Vol 2*, 16-22
- setting
 - device characteristics in, *Vol 1*, 5-12
 - printer device characteristics, *Vol 1*, 8-35
 - terminal device characteristics, *Vol 1*, 8-22
- setting up a LAT network, *Vol 1*, 5-15
- starting CONFIGURE process, *Vol 1*, 5-5
- starting InfoServer Client for OpenVMS software, *Vol 1*, 5-13
- starting LMF (License Management Facility), *Vol 1*, 5-5
- starting queues, *Vol 1*, 5-12
- starting SMISERVER process, *Vol 1*, 5-5
- starting system processes, *Vol 1*, 5-5

System startup (cont'd)

- starting the DECnet network, *Vol 1*, 5–16
- starting the queue manager, *Vol 1*, 13–3
- startup command procedures, *Vol 1*, 5–2
- submitting batch jobs, *Vol 1*, 5–14
- suppressing autoconfiguration, *Vol 1*, 5–8, 8–10
- tasks, *Vol 1*, 4–1
- VMS\$PHASES.DAT database, *Vol 1*, 5–4

System time

- See also Time
- in OpenVMS Cluster
 - daylight saving time, *Vol 1*, 6–14
 - setting, *Vol 1*, 6–13
 - standard time, *Vol 1*, 6–14
- resetting after January 1st, *Vol 1*, 6–25
- setting, *Vol 1*, 6–1 to 6–17

System tuning

- See Tuning

System users (security category)

- defining with MAXSYSGROUP parameter, *Vol 1*, 12–8
- qualifications for, *Vol 1*, 12–8

System version

- registering images with dependencies on, *Vol 1*, 5–23

System volumes

- definition, *Vol 1*, 9–14

Systemwide logical names, *Vol 1*, 5–8

SYSTEST account

- initial modification, *Vol 1*, 7–9
- in UAF, *Vol 1*, 7–7
- logging in
 - for UETP, *Vol 2*, 19–2, 19–4
- privileges required for UETP, *Vol 2*, 19–23
- quotas required to run UETP, *Vol 2*, 19–23

SYSTEST directory

- creating for UETP, *Vol 2*, 19–6
- function during UETP, *Vol 2*, 19–4

SYSTEST_CLIG account

- in UAF, *Vol 1*, 7–7
- reenabling for UETP, *Vol 2*, 19–10
- requirements for UETP, *Vol 2*, 19–10, 19–38

SYSUAF.DAT files

- definition, *Vol 1*, 7–2
- moving to reduce system disk I/O, *Vol 2*, 17–7

SYSUAFALT.DAT files, *Vol 1*, 4–11

SYSUAF logical name, *Vol 1*, 4–11

- defining during system startup, *Vol 1*, 5–8
- defining to reduce system disk I/O, *Vol 2*, 17–7

T

Tailoring a system disk

- with VMSTAILOR and DECW\$TAILOR, *Vol 1*, 5–1

Tailoring utilities (VMSTAILOR and DECW\$TAILOR), *Vol 1*, 5–1

Tape cartridge drives

- preparing for UETP, *Vol 2*, 19–7

Tape commands

- DISMOUNT, *Vol 1*, 9–59
- MOUNT, *Vol 1*, 9–25

Tape densities

- supported on Alpha, *Vol 1*, 9–13

Tape files

- See also Tape file system
- accessing at file level, *Vol 1*, 10–28
- accessing for read and write operations, *Vol 1*, 10–29
- append operation, *Vol 1*, 10–33
- closing after opening for read access, *Vol 1*, 10–32
- closing after opening for write access, *Vol 1*, 10–33
- copying, *Vol 1*, 10–37
- definition, *Vol 1*, 9–12
- locating for read and write access, *Vol 1*, 10–30
- modifying characteristics, *Vol 1*, 10–24
- reading, *Vol 1*, 10–32
- update operation, *Vol 1*, 10–33
- writing to, *Vol 1*, 10–33

Tape file system

- checking
 - continuation volume, *Vol 1*, 9–56
 - expiration date field, *Vol 1*, 10–33
- locating files, *Vol 1*, 10–30
- overwriting existing information, *Vol 1*, 10–33
- protection on, *Vol 1*, 9–26
- writing files to tape volume, *Vol 1*, 10–33

Tape mass storage control protocol

- See TMSCP

Tapes

- See also Tape volumes
- accessing file level, *Vol 1*, 10–29
- allocating drives, *Vol 1*, 9–15, 10–38
- ANSI labels, *Vol 1*, 9–26
- basic concepts of, *Vol 1*, 9–10
- blocks, *Vol 1*, 9–11
- bpi, *Vol 1*, 9–11
- commands, *Vol 1*, 9–34
- copying files from, *Vol 1*, 10–37
- deallocating drives, *Vol 1*, 9–16
- density
 - multiple (on Alpha), *Vol 1*, 9–12
- dismounting, *Vol 1*, 11–17
- DOS-11, *Vol 1*, 10–40
- enabling write cache, *Vol 1*, 9–32
- file names, *Vol 1*, 10–31
- file protection
 - See Protection
- files
 - See Tape files

Tapes (cont'd)

- file system, *Vol 1, 9-12*
 - getting information about, *Vol 1, 8-38*
 - initializing, *Vol 1, 11-14*
 - IRG (interrecord gap), *Vol 1, 9-11*
 - labeling, *Vol 1, 11-22*
 - loading on drive, *Vol 1, 10-38*
 - management
 - tasks, *Vol 1, 8-38*
 - markers on, *Vol 1, 9-11*
 - modifying device characteristics, *Vol 1, 8-38, 9-56*
 - mounting, *Vol 1, 11-16*
 - MTACP process, *Vol 1, 9-11*
 - multiple tape densities (on Alpha), *Vol 1, 9-12*
 - preparing for UETP, *Vol 2, 19-2, 19-5, 19-6*
 - protection, *Vol 1, 9-21*
 - reading from, *Vol 1, 10-32*
 - record blocking, *Vol 1, 9-11*
 - advantages, *Vol 1, 9-12*
 - record size
 - specifying, *Vol 1, 9-33*
 - sequential organization of, *Vol 1, 9-11*
 - specifying block size for, *Vol 1, 9-32*
 - standard-labeled
 - mounting, *Vol 1, 9-31*
 - structure of, *Vol 1, 9-10*
 - testing with UETP, *Vol 2, 19-32, 19-34*
 - time limit, *Vol 1, 9-75*
 - UETP test image, *Vol 2, 19-35*
 - volume label
 - overriding protection, *Vol 1, 9-33*
 - volume protection
 - See Protection
 - volumes
 - See Tape volumes
 - volume sets, *Vol 1, 9-11*
 - write cache
 - enabling, *Vol 1, 9-32*
 - writing files to, *Vol 1, 10-37*
- Tape volumes, *Vol 1, 10-39*
- See also Tape files; Tapes; Volumes
- accessibility protection, *Vol 1, 9-26*
 - accessing files on, *Vol 1, 10-28, 10-32*
 - access to, *Vol 1, 10-29*
 - continuation, *Vol 1, 9-53*
 - copying files, *Vol 1, 10-35 to 10-40*
 - deallocating, *Vol 1, 10-38*
 - dismounting, *Vol 1, 10-38*
 - file-structured, *Vol 1, 9-28*
 - foreign, *Vol 1, 9-28*
 - header labels, *Vol 1, 9-32*
 - initializing, *Vol 1, 10-38*
 - mounting, *Vol 1, 9-28, 9-31, 9-34*
 - mounting volume set, *Vol 1, 9-51*
 - mounting with automatic switching disabled, *Vol 1, 9-54*

Tape volumes (cont'd)

- mount verification, *Vol 1, 9-73*
 - overriding UIC, *Vol 1, 9-32*
 - private, *Vol 1, 9-15*
 - reading attributes of header labels, *Vol 1, 10-32*
 - reading files on, *Vol 1, 10-32*
 - searching for files on, *Vol 1, 10-20*
 - specifying record size, *Vol 1, 9-33*
 - standard-labeled, *Vol 1, 10-37*
 - copying files from, *Vol 1, 10-37*
 - wildcard characters supported, *Vol 1, 10-31*
 - write-enabling, *Vol 1, 9-35*
 - write-locking, *Vol 1, 9-51*
 - write rings, *Vol 1, 9-51*
 - writing files to, *Vol 1, 10-37*
 - writing to files on, *Vol 1, 10-33*
- TAPE_MVTIMEOUT system parameter, *Vol 1, 9-75, 9-77*
- TCP/IP (Transmission Control Protocol/Internet Protocol), *Vol 2, 23-5*
- See Compaq TCP/IP Services for OpenVMS
- TCPIP\$CONFIG
 - configuring Compaq TCP/IP Services for OpenVMS, *Vol 2, 23-6, 23-7*
- TDF (time differential factor), *Vol 1, 6-2*
 - map for determining, *Vol 1, 6-3*
 - tables of, *Vol 2, B-1*
- Template files
 - for site-specific startup, *Vol 1, 5-3*
- Temporary working directory
 - specifying alternate working device for, *Vol 1, 3-10*
- Terminal queues, *Vol 1, 14-3*
- Terminals
 - console, *Vol 1, 2-22*
 - controlling access through system passwords, *Vol 1, 12-3*
 - determining physical type, *Vol 1, 8-24*
 - disconnecting without terminating a process, *Vol 1, 8-22*
 - documenting owners of, *Vol 1, 8-22*
 - for security alarms, *Vol 2, 20-30*
 - keeping sessions on multiple, *Vol 1, 8-22*
 - LAT, *Vol 1, 8-24*
 - determining characteristics of, *Vol 1, 8-24*
 - disconnecting, *Vol 1, 8-23*
 - managing
 - tasks for, *Vol 1, 8-21*
 - operator
 - See Operator terminals
 - preparing for UETP, *Vol 2, 19-3, 19-5, 19-8, 19-14*
 - remote, *Vol 1, 8-23*
 - SET TERMINAL/INQUIRE command, *Vol 1, 8-24*
 - setting characteristics, *Vol 1, 8-21*

- Terminals
 - setting characteristics (cont'd)
 - default values, *Vol 1*, 8–21
 - in system startup, *Vol 1*, 5–12, 8–22
 - simulating users for UETP, *Vol 2*, 19–35
 - testing with UETP, *Vol 2*, 19–32, 19–34
 - UETP output, *Vol 2*, 19–34
 - UETP test image, *Vol 2*, 19–35
 - virtual
 - See Virtual terminals
- Terminal servers, *Vol 2*, 26–5
 - definition, *Vol 2*, 26–1
 - on OpenVMS system, *Vol 2*, 26–14
- Time
 - See also TDF (time differential factor)
 - modifying system time, *Vol 1*, 6–24
 - resetting after January 1st, *Vol 1*, 6–25
 - updating in a cluster, *Vol 1*, 6–24
- Time differential factor
 - See TDF
- Timed_release job status, *Vol 1*, 14–66
- Time formats, *Vol 1*, 6–17
 - predefined, *Vol 1*, 6–21
 - specifying, *Vol 1*, 6–19, 6–20
- Timeout periods
 - in SYSMAN, *Vol 1*, 2–18
 - mount verification OPCOM message, *Vol 1*, 9–75
- Timer queue entry limit, *Vol 1*, 7–46
- Timestamp functions, *Vol 1*, 10–18
- Time zone rules, *Vol 1*, 6–3
- TLZ04 tape drives
 - supported by UETP, *Vol 2*, 19–7
- TMSCP (tape mass storage control protocol), *Vol 2*, 22–2
- TP_SERVER
 - See DECdtm
- TP_SERVER process
 - stopping permanently, *Vol 2*, 27–23
 - stopping to dismount a disk, *Vol 2*, 27–17
- TQELM process limit, *Vol 1*, 7–46
- Track
 - definition, *Vol 2*, A–2
- Trailer labels
 - on tape files, *Vol 1*, 9–12
- Trailer pages, *Vol 1*, 14–39
 - file, *Vol 1*, 14–44
 - job, *Vol 1*, 14–44
- Transaction group
 - defined, *Vol 2*, 27–5
 - example, *Vol 2*, 27–5
- Transactions
 - DECdtm
 - monitoring, *Vol 2*, 27–8
- Translation modes
 - card reader, *Vol 1*, 8–40
- Transmission speed
 - setting for terminals, *Vol 1*, 8–21, 8–22
- Transports
 - LASTport, *Vol 2*, 25–5
 - local area (LAT), *Vol 2*, 26–1
 - specifying in Mail, *Vol 1*, 5–34
 - TCP/IP, *Vol 2*, 20–45
- Trigger boot
 - MOP downline load service, *Vol 2*, 24–32
- Troubleshooting
 - adding or deleting a device control library
 - module, *Vol 1*, 14–79
 - autostart queues, *Vol 1*, 14–76
 - booting problems, *Vol 1*, 4–15, 4–18
 - general printer problems, *Vol 1*, 14–72
 - holding jobs, *Vol 1*, 14–73
 - if a device is not recognized by the system, *Vol 1*, 8–8
 - jobs that will not execute, *Vol 1*, 14–73
 - jobs with characteristic mismatch, *Vol 1*, 14–75
 - OPCOM failure, *Vol 1*, 2–23
 - pending jobs, *Vol 1*, 14–73
 - print jobs with stock mismatch, *Vol 1*, 14–75
 - problems deleting a queue, form, or
 - characteristic, *Vol 1*, 14–76
 - queue manager, *Vol 1*, 13–16
 - queue problems, *Vol 1*, 14–72
 - stalled output queue, *Vol 1*, 14–76
 - startup problems, *Vol 1*, 4–16
 - system dump file for, *Vol 2*, 16–19
 - system failure, *Vol 2*, 16–19
 - system hang, *Vol 2*, 16–24
 - system startup problems, *Vol 1*, 4–17
 - UETP, *Vol 2*, 19–21
- TT2\$M_DISCONNECT characteristic
 - enabling, *Vol 1*, 8–23
 - relationship to TTY_DEFCHAR2 system
 - parameter, *Vol 1*, 8–23
 - setting up virtual terminals, *Vol 1*, 8–23
- TTDRIVER device driver
 - loading, *Vol 1*, 8–22
- TTY_DEFCHAR2 system parameter, *Vol 1*, 8–21
 - relationship to TT2Y\$M_DISCONNECT
 - characteristic, *Vol 1*, 8–23
 - setting up virtual terminals, *Vol 1*, 8–23
- TTY_DEFCHAR system parameter, *Vol 1*, 8–21
- Tuning
 - See also Performance, managing
 - definition, *Vol 2*, 17–4
 - evaluating success, *Vol 2*, 17–6
 - hardware capacity, *Vol 2*, 17–5
 - minimizing with AUTOGEN feedback, *Vol 2*, 15–10
 - predicting when required, *Vol 2*, 17–5
 - with AUTOGEN, *Vol 2*, 15–5
- TYPE command
 - tape, *Vol 1*, 10–32

U

- UAFALTERNATE system parameter, *Vol 1*, 4–11, 4–12
- UAFs (user authorization files)
 - booting with alternate, *Vol 1*, 4–11
 - checking quotas for software installation, *Vol 1*, 3–3
 - description of, *Vol 1*, 7–2
 - general maintenance, *Vol 1*, 7–6
 - initial contents, *Vol 1*, 7–6
 - initial modification, *Vol 1*, 7–9
 - listing records in, *Vol 1*, 7–24
 - logical name defining location, *Vol 1*, 5–8
 - login check, *Vol 1*, 7–5
 - modifying
 - user record, *Vol 1*, 7–23
 - network proxy, *Vol 1*, 7–37
 - preventing applications from accessing ODS-5 volumes, *Vol 1*, 10–14
 - records
 - creating multiple default, *Vol 1*, 7–24
 - resource limits for VAX and Alpha, *Vol 1*, 7–42
 - returning to the default, *Vol 1*, 4–12
 - SYSMAN checks, *Vol 1*, 2–16
 - SYSUAF.DAT, *Vol 1*, 7–2
 - user priorities, *Vol 1*, 7–34
- UETCONT00.DAT file, *Vol 2*, 19–32
 - creation of, *Vol 2*, 19–32
- UETINIDEV.DAT file, *Vol 2*, 19–31, 19–33
 - creation of, *Vol 2*, 19–32
 - format, *Vol 2*, 19–33
- UETININET.DAT file, *Vol 2*, 19–37
- UETINIT00.EXE image, *Vol 2*, 19–31
- UETINIT01.EXE image, *Vol 2*, 19–21, 19–31
- UETLOAD00.DAT file, *Vol 2*, 19–35
- UETNETS00.EXE file, *Vol 2*, 19–37
- UETP\$NODE_ADDRESS logical name, *Vol 2*, 19–12
- UETP (User Environment Test Package)
 - aborting execution, *Vol 2*, 19–15
 - compact disc drives supported by, *Vol 2*, 19–7
 - DECnet installation defaults, *Vol 2*, 19–31
 - description, *Vol 2*, 19–1
 - displaying tests as they run, *Vol 2*, 19–18
 - initialization phase, *Vol 2*, 19–31
 - interpreting output, *Vol 2*, 19–17
 - master command procedure, *Vol 2*, 19–31
 - normal completion, *Vol 2*, 19–15
 - optical disk drives supported by, *Vol 2*, 19–7
 - organization, *Vol 2*, 19–31
 - required privileges, *Vol 2*, 19–23
 - required quotas, *Vol 2*, 19–23
 - requirements for small disk systems, *Vol 2*, 19–11
 - running all phases, *Vol 2*, 19–3
 - running individual phase, *Vol 2*, 19–13
 - UETP (User Environment Test Package) (cont'd)
 - running multiple passes, *Vol 2*, 19–14, 19–20
 - starting, *Vol 2*, 19–13
 - system resource requirements, *Vol 2*, 19–2, 19–4
 - terminal and line printer output, *Vol 2*, 19–34
 - testing vector processors, *Vol 2*, 19–12
 - TLZ04 tape drive, *Vol 2*, 19–7
 - typical failures reported by, *Vol 2*, 19–21
 - VAX Vector Instruction, *Vol 2*, 19–12
- UETP.COM procedure, *Vol 2*, 19–31
 - terminating, *Vol 2*, 19–15
- UETP.LOG file, *Vol 2*, 19–15, 19–20, 19–29, 19–36
- UETPHAS00.EXE program, *Vol 2*, 19–31, 19–32
- UETUNAS00.EXE UETP test image, *Vol 2*, 19–20
- UFDs (user file directories)
 - included in MFD, *Vol 2*, A–4
 - UIC associated with, *Vol 2*, A–4
- UICs (user identification codes), *Vol 1*, 7–14
 - default protection
 - changing, *Vol 1*, 10–22
 - directory protection, *Vol 1*, 10–26
 - disk usage kept in quota file, *Vol 2*, A–9
 - for UETP, *Vol 2*, 19–6
 - identifiers, *Vol 1*, 12–11
 - interpreting, *Vol 1*, 12–7
 - member number, *Vol 1*, 7–19
 - overriding for tape volumes, *Vol 1*, 9–32
 - protection
 - of queues, *Vol 1*, 14–23
 - public volumes, *Vol 1*, 9–22
 - [0,0], *Vol 1*, 9–64
- Unicode, *Vol 2*, D–11
- Update procedures
 - See Mandatory updates
- Upgrade procedures
 - and system version dependent applications, *Vol 1*, 5–24
 - using with existing OpenVMS, *Vol 1*, 3–2
- Usage count
 - DIRECTORY/SIZE command, *Vol 1*, 9–63
 - DISKQUOTA display, *Vol 1*, 9–63
- USE command
 - in conversational boot, *Vol 1*, 4–8
 - in SYSGEN, *Vol 2*, 15–32, 15–34
- User accounts
 - changing quotas or privileges, *Vol 1*, 7–23
 - deleting, *Vol 1*, 7–25
 - disabling, *Vol 1*, 7–27
 - listing records of, *Vol 1*, 7–24
 - maintaining, *Vol 1*, 7–24
 - modifying, *Vol 1*, 7–23
 - restricting use, *Vol 1*, 7–27
 - setting up, *Vol 1*, 7–6
- User authorization, *Vol 1*, 7–5
- User authorization files
 - See UAFs

User disks
 preparing for UETP, *Vol 2*, 19–2, 19–6
 space requirements for UETP, *Vol 2*, 19–5
 testing with UETP, *Vol 2*, 19–34
 UETP error message, *Vol 2*, 19–24
 UETP test image, *Vol 2*, 19–35

User Environment Test Package
 See UETP

User file directories
 See UFDs

User files
 on public volumes, *Vol 1*, 9–14
 placement, *Vol 1*, 9–15

User Identification Code
 See UICs

User loads
 defined for UETP DECnet test, *Vol 2*, 19–37
 defining for the UETP load test, *Vol 2*, 19–14
 equation used to determine for UETP load test,
Vol 2, 19–18, 19–19

User Mail profile, *Vol 1*, 7–41

User mode
 logical names, *Vol 2*, 17–23

User names
 as identifiers, *Vol 1*, 12–11
 entering when logging in, *Vol 1*, 2–10

User resources, *Vol 1*, 7–41

Users
 limiting the number of interactive, *Vol 1*, 5–16;
Vol 2, 17–4
 protection code categories, *Vol 1*, 12–7
 restricting login hours for, *Vol 2*, 17–4
 security categories of, *Vol 1*, 12–7 to 12–8
 sending messages to with OPCOM, *Vol 1*, 2–23
 sending requests to an operator, *Vol 1*, 2–26
 validation of, *Vol 1*, 2–16

User-specified job retention
 changing, *Vol 1*, 14–29

UTC (Coordinated Universal Time), *Vol 1*, 6–1

UTC service
 See Coordinated Universal Time (UTC)

Utilities for configuring OpenVMS Registry, *Vol 2*,
 D–1

V

Validation of users, *Vol 1*, 2–16

VAXcluster environments
 defining the number of nodes for AUTOGEN,
Vol 2, 15–21

VAXcluster systems
 adjusting quorum after shutting down a node,
Vol 1, 4–24
 shutting down an entire VAXcluster, *Vol 1*,
 4–24

VAX Vector Instruction Emulation Facility
 (VVIEF), *Vol 2*, 19–12
 definition, *Vol 2*, 28–5
 determining presence of, *Vol 2*, 28–9, 28–10
 loading, *Vol 2*, 28–10
 unloading, *Vol 2*, 28–10

VAXVMSSYS.PAR file, *Vol 1*, 4–2; *Vol 2*, 15–3
 initializing parameters at boot time, *Vol 2*,
 15–36

VBN_CACHE_S system parameter, *Vol 2*, 18–3,
 18–13

VCC_FLAGS system parameter, *Vol 2*, 18–3,
 18–13

VCC_MAXSIZE system parameter, *Vol 2*, 18–11,
 18–14

VCC_MAX_CACHE system parameter, *Vol 2*,
 18–5

VCC_MAX_IO_SIZE system parameter, *Vol 2*,
 18–6

VCC_PTES system parameter, *Vol 2*, 18–11

VCC_READAHEAD system parameter, *Vol 2*,
 18–8

VCRs (vector count registers), *Vol 2*, 28–4

Vector capability
 determining availability within a system, *Vol 2*,
 28–9
 placing an ACL on, *Vol 2*, 28–8

Vector-capable, *Vol 2*, 28–4

Vector consumer
 determining the identity of, *Vol 2*, 28–9
 managing, *Vol 2*, 28–6
 marginal, *Vol 2*, 28–8
 obtaining information about, *Vol 2*, 28–8

Vector context switch
 obtaining information about, *Vol 2*, 28–9

Vector CPU time
 obtaining information regarding process, *Vol 2*,
 28–9

Vector-present processors, *Vol 2*, 28–4
 adding to system, *Vol 2*, 28–6
 identifying, *Vol 2*, 28–9
 removing from system, *Vol 2*, 28–6
 when unavailable, *Vol 2*, 28–6

Vector processing, *Vol 2*, 28–4 to 28–9
 configuring system, *Vol 2*, 28–6
 definition, *Vol 2*, 28–4
 establishing batch queues for, *Vol 2*, 28–7
 managing, *Vol 2*, 28–5
 obtaining information about, *Vol 2*, 28–8
 obtaining number of vector processors, *Vol 2*,
 28–9
 resource requirements, *Vol 2*, 28–7
 system performance, *Vol 2*, 28–4
 tasks for managing, *Vol 2*, 28–5
 tuning system, *Vol 2*, 28–7
 VAX support for, *Vol 2*, 28–4

- Vector registers, *Vol 2*, 28–4
- Vectors
 - definition, *Vol 2*, 28–4
- VECTOR_MARGIN system parameter, *Vol 2*, 28–8
- VECTOR_PROC system parameter, *Vol 2*, 28–5
- Verification
 - mount
 - See Mount verification
 - startup, *Vol 1*, 4–17
 - turning on during system startup, *Vol 1*, 4–16
- Version dependencies
 - registering images with, *Vol 1*, 5–23
- Version limits
 - setting on files, *Vol 1*, 9–68
- Version numbers, *Vol 1*, 10–30
- Virtual devices
 - mounting during system startup, *Vol 2*, 25–16
- Virtual I/O Cache, *Vol 2*, 18–1
- Virtual I/O caches, *Vol 2*, 18–10 to 18–14
 - adjusting size, *Vol 2*, 18–14
 - caching policies, *Vol 2*, 18–10
 - disabling, *Vol 2*, 18–13
 - enabling, *Vol 2*, 18–13
 - memory allocation, *Vol 2*, 18–13
 - memory reclamation, *Vol 2*, 18–13
 - objects, *Vol 2*, 18–10
 - OpenVMS Cluster configurations, *Vol 2*, 18–14
 - parameters
 - VBN_CACHE_S, *Vol 2*, 18–13
 - VCC_FLAGS, *Vol 2*, 18–13
 - VCC_MAXSIZE, *Vol 2*, 18–14
 - statistics, *Vol 2*, 18–12
- VIRTUALPAGECNT system parameter
 - optimizing batch queues for Sort/Merge utility, *Vol 1*, 14–35
- Virtual terminals
 - connecting, *Vol 1*, 8–22
 - determining physical terminal type, *Vol 1*, 8–24
 - enabling, *Vol 1*, 8–22
 - purpose of, *Vol 1*, 8–22
 - TT2\$M_DISCONNECT characteristic, *Vol 1*, 8–22
 - TTY_DEFCHAR2 system parameter, *Vol 1*, 8–23
- VLRs (vector length registers), *Vol 2*, 28–4
- VMB.EXE file, *Vol 1*, 4–20
 - role in boot process, *Vol 1*, 4–2
- VMRs (vector mask registers), *Vol 2*, 28–4
- VMS\$LAYERED.DAT file, *Vol 1*, 5–19
 - function in startup procedure, *Vol 1*, 5–19
- VMS\$PHASES.DAT file
 - in startup procedure, *Vol 1*, 5–4, 5–18
- VMS\$VMS.DAT file
 - in startup procedure, *Vol 1*, 5–19
- VMSINSTAL.COM command procedure
 - See also Installation procedures; Installing software
 - Alternate System Root option, *Vol 1*, 3–14
 - restriction, *Vol 1*, 3–14
 - Alternate Working Device option, *Vol 1*, 3–10
 - answer file, *Vol 1*, 3–10
 - Autoanswer option, *Vol 1*, 3–10
 - BACKUP qualifiers, *Vol 1*, 3–12
 - command line syntax, *Vol 1*, 3–5
 - completing installation, *Vol 1*, 3–9
 - correcting problems, *Vol 1*, 3–5
 - creating new answer file, *Vol 1*, 3–10
 - destination parameter, *Vol 1*, 3–9
 - File Log option, *Vol 1*, 3–13
 - Get Save Set option, *Vol 1*, 3–11, 3–12
 - getting help in, *Vol 1*, 3–5
 - installing layered products, *Vol 1*, 3–2
 - options, *Vol 1*, 3–10
 - option list parameter, *Vol 1*, 3–8
 - specifying, *Vol 1*, 3–8
 - table of, *Vol 1*, 3–8
 - preparing to use, *Vol 1*, 3–3
 - product list parameter, *Vol 1*, 3–5
 - product save-set format, *Vol 1*, 3–12
 - Release Notes option, *Vol 1*, 3–13
 - saving answers, *Vol 1*, 3–10
 - source parameter, *Vol 1*, 3–7
 - starting, *Vol 1*, 3–5
 - system failure
 - conditions for, *Vol 1*, 3–9
 - system shutdown following, *Vol 1*, 3–9
 - temporary working directory
 - specifying location of, *Vol 1*, 3–11
- VMSKITBLD.COM command procedure
 - ADD option, *Vol 1*, 2–32
 - BUILD option, *Vol 1*, 2–28
 - compared to CLUSTER_CONFIG.COM
 - command procedure, *Vol 1*, 2–32
 - completing a system disk created with the BUILD option, *Vol 1*, 2–30
 - configuring a system root added with, *Vol 1*, 2–34
 - copying system files from the system disk, *Vol 1*, 2–27
 - COPY option, *Vol 1*, 2–31
 - options, *Vol 1*, 2–27
 - reliance on .TEMPLATE version of site-specific command procedures, *Vol 1*, 5–3
- VMSMAIL_PROFILE.DATA file, *Vol 1*, 7–41
 - moving to reduce system disk I/O, *Vol 2*, 17–7
- VMSMAIL_PROFILE logical name
 - defining during system startup, *Vol 1*, 5–8
 - defining to reduce system disk I/O, *Vol 2*, 17–7
- VMSTAILOR
 - See Tailoring utilities

- VOLPRO privilege, *Vol 1*, 9–21
- VOLSET.SYS file
 - See Volume set list file
- Volume
 - definition, *Vol 1*, 9–2
- Volume identifier field, *Vol 1*, 9–53
- Volume integrity, *Vol 1*, 9–73
- Volume labels
 - assigning to devices, *Vol 1*, 5–11
 - changing, *Vol 1*, 3–30
 - definition, *Vol 1*, 11–59
 - format, *Vol 1*, 9–31
 - specifying for magnetic tape, *Vol 1*, 11–15
 - used with BACKUP command, *Vol 1*, 11–59
- Volume protection
 - See Protection, volumes
- Volume Protection Override (VOLPRO)
 - See VOLPRO privilege
- Volumes
 - See also Disk volumes; Private volumes; Public volumes; Tape volumes; Volume sets
 - adding to an existing disk set, *Vol 1*, 9–47
 - advantages of using separate, *Vol 1*, 9–43
 - availability
 - OPCOM message, *Vol 1*, 9–73
 - binding into volume set, *Vol 1*, 9–29
 - canceling mount verification, *Vol 1*, 9–77
 - continuation, *Vol 1*, 9–56
 - controlling cache size, *Vol 1*, 9–73
 - disabling mount messages, *Vol 1*, 9–31
 - disk
 - See Disk volumes
 - dismounting, *Vol 1*, 9–57, 11–17
 - restriction, *Vol 2*, 17–26
 - foreign
 - copying files to and from, *Vol 1*, 10–40
 - mounting, *Vol 1*, 9–28, 10–28
 - group, *Vol 1*, 9–14
 - initializing, *Vol 1*, 9–16, 11–14
 - mounting, *Vol 1*, 9–35, 11–16
 - See also MOUNT command
 - continuation tape, *Vol 1*, 9–54
 - if device is unavailable, *Vol 1*, 9–35
 - operator functions, *Vol 1*, 9–25
 - public, *Vol 1*, 9–27
 - steps, *Vol 1*, 9–34
 - mount verification aborted
 - OPCOM message, *Vol 1*, 9–79
 - mount verification timeout, *Vol 1*, 9–75
 - OPCOM message, *Vol 1*, 9–75
 - operator-assisted mount, *Vol 1*, 9–35
 - private
 - See Private volumes
 - protection, *Vol 1*, 9–21
 - public
 - See Public volumes

- Volumes (cont'd)
 - rebuilding, *Vol 1*, 9–64
 - determining if needed, *Vol 1*, 8–5
 - recovering from errors, *Vol 1*, 9–74
 - sharing, *Vol 1*, 9–31
 - substituting, *Vol 1*, 9–35
 - system
 - See System volumes
 - tape
 - See Tape volumes
- Volume security profile
 - reserved file, *Vol 2*, A–9
 - SECURITY.SYS, *Vol 2*, A–9
- Volume set list file
 - definition, *Vol 2*, A–9
 - reserved file, *Vol 2*, A–9
 - used by Analyze/Disk_Structure utility, *Vol 2*, A–9
 - VOLSET.SYS, *Vol 2*, A–9
- Volume sets
 - backing up, *Vol 1*, 11–43
 - CD-ROM
 - partially mounted ISO 9660, *Vol 1*, 9–48
 - characteristics, *Vol 1*, 9–43
 - creating, *Vol 1*, 9–29, 9–43
 - definition, *Vol 1*, 9–3, 9–43
 - disk, *Vol 1*, 9–44 to 9–47
 - accessing, *Vol 1*, 9–44, 9–45
 - adding volumes, *Vol 1*, 9–47
 - assigning name to, *Vol 1*, 9–44
 - creating, *Vol 1*, 9–44
 - from existing volumes, *Vol 1*, 9–46
 - from new volumes, *Vol 1*, 9–45
 - shadowed, *Vol 1*, 9–46
 - creating files on, *Vol 1*, 9–44
 - directory structure, *Vol 1*, 9–44
 - index file, *Vol 1*, 9–45
 - naming, *Vol 1*, 9–44
 - information stored in VOLSET.SYS, *Vol 2*, A–9
 - mounting, *Vol 1*, 9–44
 - See also MOUNT command
 - privileges to access, *Vol 1*, 9–45
 - processing continuation volumes, *Vol 1*, 9–51
 - restoring, *Vol 1*, 11–50
 - restriction for system disk, *Vol 1*, 9–44
 - root volume, *Vol 1*, 9–44
 - tape, *Vol 1*, 9–11, 9–51 to 9–54
 - automatic volume switching, *Vol 1*, 9–53
 - continuation volumes, *Vol 1*, 9–52
 - creating, *Vol 1*, 9–51
 - mounting, *Vol 1*, 9–51
 - mounting continuation volumes, *Vol 1*, 9–54
 - mounting with automatic switching disabled, *Vol 1*, 9–54

- Volume shadowing, *Vol 1*, 11–43, 11–50
 - minimerge (assisted merge), *Vol 1*, 11–45
 - mounting disks in host-based shadow set, *Vol 1*, 11–45
 - RAID, *Vol 1*, 11–43
- Volume switching
 - automatic, *Vol 1*, 9–53
- VTA0 device
 - connecting, *Vol 1*, 8–22
- VVIEF
 - See VAX Vector Instruction Emulation Facility
- VVIEF\$DEINSTAL.COM command procedure, *Vol 2*, 28–10
- VVIEF\$INSTAL.COM command procedure, *Vol 2*, 28–5, 28–10

W

- Welcome messages
 - defining, *Vol 1*, 5–14
 - displaying, *Vol 1*, 5–15
 - login, *Vol 1*, 2–10
- Wildcard characters
 - asterisk (*)
 - using with tape volumes, *Vol 1*, 10–31
 - in file names, *Vol 1*, 10–31
 - with OpenVMS extended names, *Vol 1*, 10–31
 - with standard names, *Vol 1*, 10–31
 - with tape volumes, *Vol 1*, 10–31
- Working directory
 - temporary
 - in VMSINSTAL.COM command procedure, *Vol 1*, 3–11
- Working sets
 - adjusting for vectorized applications, *Vol 2*, 28–7
 - default size, *Vol 1*, 7–46
 - extent, *Vol 1*, 7–46
 - limits and quotas
 - choosing a value for batch queues, *Vol 1*, 14–34, 14–35
 - specifying for batch queues, *Vol 1*, 14–22, 14–32
 - specifying for output queues, *Vol 1*, 14–22
 - quota, *Vol 1*, 7–47
- Work load
 - adjusting system parameters for, *Vol 2*, 17–5
 - distributing, *Vol 2*, 17–3
 - designing efficient applications, *Vol 2*, 17–4
 - restricting login hours, *Vol 2*, 17–4
 - knowing, *Vol 2*, 17–2
 - monitoring, *Vol 2*, 17–2
 - tools used for, *Vol 2*, 17–2
 - strategies for managing, *Vol 2*, 17–3
- Workstations
 - backing up, *Vol 1*, 11–39
 - managing queues on, *Vol 1*, 14–2
 - OPCOM behavior on, *Vol 1*, 2–22

- Workstations (cont'd)
 - OPCOM startup on, *Vol 1*, 2–22
 - printer queue configuration, *Vol 1*, 14–8
 - setting up media on, *Vol 1*, 9–17
 - starting SMISERVER process, *Vol 1*, 2–12
- World users (security category), *Vol 1*, 12–8
- Writable images, *Vol 2*, 17–19
- Write access
 - See also Access types, write
 - for disk directory files, *Vol 1*, 10–26
 - for disk files, *Vol 1*, 10–21
 - granting through protection codes, *Vol 1*, 12–8
- Writeboot utility (WRITEBOOT), *Vol 1*, 4–19
 - error messages, *Vol 1*, 4–21
- Write cache
 - enabling for tape device, *Vol 1*, 9–32
- Write-enabling a tape, *Vol 1*, 9–35
- Write-lock errors
 - recovering from, *Vol 1*, 9–76
- Write-locking
 - disk volumes, *Vol 1*, 9–74
 - tape volumes, *Vol 1*, 9–51
- Write operation
 - See Access types, write
- Write rings
 - on tape volumes, *Vol 1*, 9–51
- Write-through caching, *Vol 2*, 18–2
- WSDEFAULT process limit
 - choosing a value for batch queues, *Vol 1*, 14–34
 - specifying a value for batch queues, *Vol 1*, 14–22, 14–32
 - specifying a value for output queues, *Vol 1*, 14–22
 - specifying values, *Vol 1*, 14–22 to 14–34
- WSdef process limit, *Vol 1*, 7–46
- WSextent process limit
 - specifying a value for batch queues, *Vol 1*, 14–22
 - specifying a value for output queues, *Vol 1*, 14–22
- WSEXTENT process limit, *Vol 1*, 7–46
 - choosing a value for batch queues, *Vol 1*, 14–34
 - for efficient sorting, *Vol 1*, 14–35
 - specifying a value for batch queues, *Vol 1*, 14–32
 - value for efficient backups, *Vol 1*, 11–11
- WSMAX system parameter, *Vol 1*, 14–34; *Vol 2*, 28–7
- WSquo process limit, *Vol 1*, 7–47
- WSQUOTA process limit
 - choosing a value for batch queues, *Vol 1*, 14–34
 - specifying a value for batch queues, *Vol 1*, 14–23, 14–32
 - specifying a value for output queues, *Vol 1*, 14–23
 - value for efficient backups, *Vol 1*, 11–11

X

XA Gateway, *Vol 2*, 27–24

XARs (extended attribute records)

keyword with MOUNT/PROTECTION

command, *Vol 1*, 9–30

protection fields, *Vol 1*, 9–24

mount option, *Vol 1*, 9–24

XDELTA debugger

See Delta/XDelta Debugger (DELTA/XDELTA)

X terminal client, *Vol 2*, 25–5