
HP OpenVMS Alpha Version 7.3–2 Release Notes

Order Number: AA-RV8YA-TE

September 2003

This manual describes changes to the software; installation, upgrade, and compatibility information; new and existing software problems and restrictions; and software and documentation corrections.

Revision/Update Information:	This is a new manual.
Software Version:	OpenVMS Alpha Version 7.3-2

**Hewlett-Packard Company
Palo Alto, California**

© Copyright 2003 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Proprietary computer software. Valid license from HP 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.

Intel is a trademark of Intel Corporation in the U.S. and other countries and is used under license.

Java™ is a US trademark of Sun Microsystems, Inc.

Oracle® is a US registered trademark of Oracle Corporation, Redwood City, California.

OSF and Motif are trademarks of The Open Group in the US and other countries.

UNIX® is a registered trademark of The Open Group.

Microsoft®, Windows®, Windows NT®, and MS Windows® are US registered trademarks of Microsoft Corporation.

X/Open® is a registered trademark, and the X device is a trademark of X/Open Company Ltd. in the UK and other countries.

ZK6668

The HP OpenVMS documentation set is available on CD-ROM.

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

Contents

Preface	ix
1 OpenVMS Alpha Software Installation and Upgrade Release Notes	
1.1 HP Software Technical Support Policy	1-1
1.2 General Application Compatibility Statement	1-2
1.3 Firmware for OpenVMS Alpha Version 7.3-2	1-3
1.4 Upgrade Paths	1-3
1.4.1 Cluster Rolling Upgrades: Correction to <i>HP OpenVMS Alpha Version 7.3-2 Upgrade and Installation Manual</i>	1-4
1.5 Networking Options	1-4
1.6 Obtaining Remedial Kits	1-5
1.7 System Startup Failure in a Cluster	1-5
1.8 SHADOW_MAX_UNIT Default Setting and Memory Usage	1-5
1.9 Associated Products Affecting OpenVMS Installation and Upgrades	1-6
1.9.1 Some Product Kits Ship in Compressed Format	1-6
1.9.2 ACMS Kits: Problem with File Deletions	1-7
1.9.3 COM for OpenVMS Upgrade in Clusters Running OpenVMS Version 7.2-1 or 7.2-1H1	1-7
1.9.4 DECEvent Version 3.1 or Higher Required for Certain Systems	1-7
1.9.5 DECnet-Plus and DECwindows Require New Versions	1-8
1.9.6 DECnet-Plus Installation: PCSI-I-RETAIN Messages	1-9
1.9.7 HP PATHWORKS and HP Advanced Server for OpenVMS Products	1-9
1.9.7.1 HP Advanced Server for OpenVMS	1-9
1.9.7.2 HP PATHWORKS for OpenVMS (Advanced Server)	1-10
1.9.7.3 PATHWORKS V5 for OpenVMS (LAN Manager) Not Supported	1-10
1.9.7.4 Upgrading Systems with Pre-V6.1 PATHWORKS Advanced Server	1-10
1.9.7.5 Upgrading Advanced Server Version 7.2x for OpenVMS	1-10
1.9.8 HP X.25 Version 1.6 Required Upgrade	1-11
1.9.9 Kerberos V1.0: Remove Before Upgrading	1-11
1.9.10 Registry Upgrade from OpenVMS Version 7.2-1 or 7.2-1H1	1-12
1.9.11 SSL: Installing HP SSL V1.1-A	1-14
1.9.12 System Event Analyzer (SEA) Utility	1-15

2 OpenVMS Associated Products Release Notes

2.1	Associated Product Support	2-1
2.2	ACMS: Changes in Error Reporting	2-1
2.3	BASIC: V1.5A Required to Create STARLET Library	2-2
2.4	COBOL: Changes in I/O Run-Time Diagnostics and RMS Special Registers	2-2
2.5	COM for OpenVMS	2-3
2.5.1	NTA\$LOGON /TYPE Qualifier Behavior Changes	2-3
2.5.2	Registry Access Error with Heavy Load of Applications	2-3
2.6	DECdfs Version 2.3-3 Required for OpenVMS Version 7.3-2	2-3
2.7	DECforms Web Connector Version 3.0	2-4
2.8	DEC PL/I: RTL Support for OpenVMS	2-4
2.9	FMS ECO Kit	2-4
2.10	Graphical Configuration Manager (GCM)	2-5
2.11	HP DECram	2-5
2.11.1	DECram Version 3.2 Required for OpenVMS Version 7.3-2	2-5
2.11.2	DECram Version 2.5	2-5
2.11.3	Maximum Disk Size for DECram Version 3.*	2-5
2.11.4	Earlier Versions of DECram	2-5
2.11.5	DECram Commands and Errors	2-6
2.11.6	DECram and Volume Shadowing	2-6
2.12	HP DECwindows Motif for OpenVMS	2-6
2.12.1	Supported Versions	2-6
2.12.2	Support for LAT Transport Interface Restored	2-6
2.12.3	User-Written Transports No Longer Supported	2-7
2.12.4	Available Language Variants	2-7
2.12.5	System Parameter Values Required for Installation	2-7
2.13	MultiNet Version 4.3	2-8
2.14	Pascal	2-8
2.14.1	V5.8A Required to Create STARLET Library	2-8
2.14.2	Installing HP Pascal After an Upgrade	2-8

3 General User Release Notes

3.1	OpenVMS Freeware CD-ROMs	3-1
3.2	Online Help Topic Name Changes	3-2
3.3	COPY Command Performance	3-2
3.4	ATI RADEON 7500 Graphics	3-2
3.5	HP Secure Web Browser: Increased Memory Required	3-2

4 System Management Release Notes

4.1	ACME_SERVER has Multiple Channels to RIGHTSLIST.DAT	4-1
4.2	AUTOGEN: New Behavior with NEWPARAMS.DAT Files	4-1
4.3	DECdtm Services	4-2
4.3.1	DECdtm/XA with Oracle® 8i and 9i (Alpha Only)	4-2
4.3.2	IPC-E-BCKTRNSFAIL Error Message	4-2
4.4	ECP Data Collector and Performance Analyzer V5.4D	4-3
4.5	EDIT/FDL: Fixing Recommended Bucket Size	4-3
4.6	Error Log Viewer (ELV) Utility	4-3
4.6.1	Translating Error Log Files from OpenVMS Versions 7.2 through 7.3-1	4-3
4.6.2	Using the /PAGE Qualifier with the TRANSLATE Command	4-3
4.7	Extended File Cache (XFC) Remedial Kits	4-4

4.8	External Authentication	4-4
4.8.1	SET PASSWORD Behavior Within a DECterm Terminal Session	4-4
4.8.2	No Password Expiration Notification on Workstations	4-5
4.9	INITIALIZE Command: Incorrect Messages Are Output	4-5
4.10	Lock Manager: Fast Lock Remastering and PE1	4-5
4.11	Logical Disk (LD) Utility: Error when Using RMS	4-6
4.12	MAIL Utility: Documentation Correction	4-6
4.13	OpenVMS Cluster Systems	4-6
4.13.1	Patch Kits Needed for Cluster Compatibility	4-6
4.13.2	New API Can Correct Incompatibility of Fibre Channel and SCSI Multipath with Some Third-Party Products	4-8
4.13.3	SCSI Tape Drives: MEDOFL Errors After Tape Dismount	4-9
4.13.4	CLUSTER_CONFIG.COM and Limits on Root Directory Names	4-9
4.13.5	Booting Satellites Over FDDI in a Mixed-Version Cluster	4-10
4.13.6	PEdriver Error Message Change	4-10
4.13.7	PEdriver Channels with Priority of -128 Not Used	4-11
4.13.8	Cluster Performance Reduced with CI-LAN Circuit Switching	4-11
4.13.9	Gigabit Ethernet Switch Restriction in an OpenVMS Cluster System	4-12
4.13.10	Multipath Tape Failover Restriction	4-12
4.13.11	No Automatic Failover for SCSI Multipath Medium Changers	4-13
4.14	OpenVMS Galaxy	4-13
4.14.1	OpenVMS Graphical Configuration Manager	4-13
4.14.2	Smart Array 5300 Restrictions	4-13
4.14.3	Firmware and Patch Kit Requirements	4-13
4.14.4	Shared-Memory Global Section Creation Can Return Incorrect Status	4-14
4.14.5	Galaxy on ES40: Uncompressed Dump Limitation	4-14
4.14.6	Galaxy on ES40: Turning Off Fast Path	4-14
4.15	OpenVMS Management Station	4-14
4.16	OpenVMS Registry Can Corrupt Version 2 Format Database	4-15
4.17	RMS Journaling	4-15
4.17.1	Recovery Unit Journaling Incompatible with Kernel Threads	4-15
4.17.2	Modified Journal File Creation	4-15
4.17.3	Remote Access of Recovery Unit Journaled Files in an OSI Environment	4-16
4.17.4	After-Image (AI) Journaling	4-16
4.17.5	VFC Format Sequential Files	4-16
4.18	Security: Changes to DIRECTORY Command Output	4-17
4.19	Server Management Process (SMHANDLER)	4-17
4.20	SYSGEN: Security Auditing Fixed	4-17
4.21	SYSMAN Utility	4-18
4.21.1	DUMP_PRIORITY LOAD Command	4-18
4.21.2	XP LUNs Can Fail to Autoconfigure	4-18
4.22	System Parameters	4-19
4.22.1	New System Parameters	4-19
4.22.2	Modified System Parameters	4-19
4.22.3	Obsolete System Parameters	4-20
4.22.4	SCSNODE Parameter Size Now Strictly Enforced	4-20
4.22.5	AlphaServer GS Series: NPAGERAD System Parameter Default Behavior	4-20
4.23	Terminal Fallback Facility (TFF)	4-20
4.24	UETP: Error With DEGPA Device	4-21
4.25	VMSINSTAL: NOBROADCAST Option	4-22

4.26	Volume Shadowing for OpenVMS	4-22
4.26.1	Device Name Requirement	4-22
4.26.2	Warning About Using SET SHADOW and SHOW SHADOW in DCL Command Procedures	4-22
4.26.3	Dissimilar Device Shadowing (DDS): Restriction on Adding First Dissimilar Member	4-23
4.26.4	Write Bitmaps and Dissimilar Device Shadowing (DDS) Caution	4-23
4.26.5	KZPDC (Smart Array 5300) Restrictions	4-24
4.26.6	SHOW DEVICE/BITMAP Command Might Produce SYSTEM-F-INTDIV Errors	4-25
4.26.7	SHOW DEVICE/BITMAP Requires LOG_IO Privilege	4-25
4.26.8	Changes in Shadow Set Merge Delay Computation	4-26
4.26.9	Deleting a Fibre Channel Device from a Controller Pair	4-26
4.26.10	ANALYZE/DISK/SHADOW Command Behavior	4-26
4.26.11	ANALYZE/DISK/SHADOW Command Behavior with Dissimilar Device Shadow Sets	4-27
4.26.12	Dismount of Shadow Set Member Using /MINICOPY	4-27

5 Programming Release Notes

5.1	Backup API: Journaling Callback Events Restriction	5-1
5.2	C Programs: Compiling with case=sensitive Settings	5-1
5.3	C Run-Time Library	5-1
5.3.1	Undefined Symbol GETADDRINFO_COMPAT43 Fixed	5-1
5.3.2	64-Bit Pointers Now Supported for Some Socket Routines	5-2
5.3.3	fwrite Function No Longer Writes Inappropriate Records to a File . . .	5-2
5.3.4	FILE_TYPE Macro Removed from <stdio.h> Header	5-2
5.3.5	tm_t Structure Removed from <time.h> Header	5-2
5.3.6	<ftw.h> Header File Now Compatible with Local-Time Function ftw	5-2
5.3.7	Incomplete Support for 64-bit iconv() Removed	5-3
5.3.8	dlsym Function No Longer Fails on Images Linked with Warnings . . .	5-3
5.3.9	dlsym Function Now Recognizes Lowercase Names	5-3
5.3.10	catopen Function Now Supports UNIX® Style Path Names	5-3
5.3.11	Memory Management Functions Are Now Reentrant	5-4
5.3.12	Angle Brackets in File Specs Now Handled Correctly	5-4
5.3.13	Race Condition in sleep and usleep Is Fixed	5-4
5.3.14	UNIX to OpenVMS Name Translation Errors Fixed	5-4
5.3.15	Behavior Changed for tzset with Invalid TZ	5-4
5.3.16	Attempt to Perform a pread or pwrite on an Unseekable File Is Now Properly Flagged as an Error	5-5
5.3.17	Open Files Removal Now Controlled With Feature Logical	5-5
5.3.18	fstat Function Now Checks Whether Disk Access Dates Were Enabled	5-5
5.3.19	Problem With Redefinition of Structures When Including <time.h> or <signal.h>	5-5
5.3.20	Mixed Logical Name Search Lists Now Work	5-7
5.3.21	Program Deadlock Problem is Fixed	5-8
5.3.22	inet_ntop function Now Defined as Returning "const char *"	5-8
5.3.23	exec No Longer Gets 0x35DF94 Error	5-8
5.3.24	Compiler Errors Gone After Upgrading to OpenVMS Version 7.3-1 . . .	5-8
5.3.25	Missing Entry Points for BSD 4.4 Functions Fixed	5-8
5.3.26	Correct IPv6 Symbol Names Now in UCX\$IPC_XHR.EXE	5-9
5.3.27	Translation Problem with Certain UNIX File Names Is Fixed	5-9

5.3.28	rename Function Now Handles Logical Names Correctly	5-9
5.3.29	IPv6 Programming Note	5-10
5.3.30	glob Function Not Supported with /POINTER_SIZE=LONG	5-10
5.3.31	New 64-bit Function Support: Special Considerations	5-10
5.4	Common Data Security Architecture (CDSA) Considerations	5-12
5.5	CONVERT-I-SEQ Error on CONVERT/NOSORT with Collated Key	5-13
5.6	Debugger: Previous Versions of Client/Server Interface Not Supported	5-13
5.7	Debugging Modes: Avoiding CPUSPINWAIT Bugchecks	5-14
5.8	HP COBOL Run-Time Library (RTL) for OpenVMS Alpha	5-14
5.8.1	COBOL RTL and UNSTRING	5-14
5.8.2	COBOL RTL Detects Device Full on CLOSE	5-14
5.8.3	COBOL RTL and Record-Locking Restriction	5-14
5.9	Hypersort Utility	5-15
5.9.1	CONVERT Performance with Large Files	5-15
5.9.2	Hypersort and VFC Files Restriction	5-15
5.9.3	SORT and MERGE Can Hang with Large Input Files	5-15
5.9.4	/FORMAT=RECORD_SIZE Restriction	5-15
5.9.5	Using Hypersort with Search Lists and Other Uses of Logical Names	5-15
5.9.6	Lack of Free Space for Work Files	5-16
5.9.7	Hypersort and SORT32 Performance: Working Set and Page File Quota	5-16
5.9.8	Hypersort and SORT32 Performance with Variable Length Records	5-16
5.9.9	Input Asterisk (*) Restriction	5-16
5.9.10	Work File Directories Restriction	5-16
5.10	Librarian Utility: PGFLQUOTA Should Exceed 23000	5-16
5.11	Lightweight Directory Access Protocol (LDAP) API	5-17
5.11.1	The Routine ldap_get_option Returns Error -1 When ld Is NULL	5-17
5.11.2	The Routine ber_flatten() Fails to Detect Mismatched Braces	5-17
5.12	Linker Utility	5-17
5.12.1	Linker Appears to Hang When Many Files Are Specified	5-17
5.12.2	Change in Linker Default Behavior with Library Check	5-19
5.12.3	Limit of 25 Elements on Stack	5-19
5.13	LTDRIVER: CANCEL SELECTIVE Restriction	5-19
5.14	Mail Utility: Threads Restriction for Callable Mail	5-19
5.15	POSIX Threads Library	5-19
5.15.1	C Language Compilation Header File Changes	5-20
5.15.2	New Priority Adjustment Algorithm	5-20
5.15.3	Process Dumps	5-21
5.15.4	Dynamic CPU Configuration Changes	5-21
5.15.5	Enhanced Debugging of Threaded Programs	5-21
5.15.6	Debugger Metering Function Does Not Work	5-22
5.16	Privileged Interfaces and Data Structures	5-22
5.16.1	Per-Thread Security Impacts Privileged Code and Device Drivers	5-22
5.16.2	IPL Requirement For OpenVMS Fork Thread Creation Now Enforced	5-23
5.17	RTL Library (LIB\$)	5-23
5.17.1	LIB\$FIND_IMAGE_SYMBOL: Error in the <i>OpenVMS RTL Library</i> (<i>LIB\$</i>) <i>Manual</i>	5-23
5.17.2	LIB\$FIND_IMAGE_SYMBOL Signals Warning	5-24
5.18	Screen Management (SMG\$) Documentation	5-24
5.19	SORT32 Utility	5-25

5.19.1	VFC Files Now Written Correctly	5-25
5.19.2	SORT32 May Not Delete Temporary Work Files	5-25
5.19.3	SORT/SPECIFICATION and Compound Conditions Restriction	5-25
5.19.4	SORT32 and Hypersort Performance: Working Set and Page File Quota	5-25
5.19.5	SORT32 and Hypersort Performance with Variable Length Records	5-25
5.19.6	SORT32 Work File Directories Restriction	5-26
5.20	System Services: SYS\$ACM(W) Changes	5-26
5.21	Timer Queue Entries (TQEs)	5-26

6 Hardware Release Notes

6.1	ALPHABook 1	6-1
6.2	AlphaServer 1000A	6-1
6.2.1	Bus Probe Algorithm Default	6-2
6.2.2	Installation Failure with DEFPA Adapter	6-2
6.3	AlphaServer 2100	6-2
6.3.1	Console Display	6-2
6.3.2	SCSI Controller Restriction	6-3
6.4	AlphaServer 8200/8400: FRU Table Error	6-3
6.5	AlphaServer ES47/ES80/GS1280 Systems	6-3
6.5.1	Firmware Version 6.6 Recommended	6-3
6.5.2	Setting SYSGEN Parameter PHYSICAL_MEMORY	6-4
6.5.3	RAD Support	6-4
6.5.4	License Requirements	6-4
6.5.5	TKNOVF Error on Output for Some F\$GETSYI Item Codes	6-4
6.5.6	STOP/CPU and Shutdown Behavior	6-4
6.5.7	Setting Time at MBM	6-5
6.5.8	ERLBUFFERPAGES Must Be Increased for AlphaServer GS1280 Systems	6-5
6.6	AlphaServer GS Series Systems	6-5
6.6.1	AlphaServer GS80/160/320 Systems: Device Restriction	6-5
6.6.2	OpenVMS Galaxy License Enforcement	6-5
6.6.3	License with Hard and Soft Partitions—Problem Corrected	6-6
6.6.4	Removing Workarounds	6-6
6.6.5	Installing Licenses	6-7
6.6.6	AlphaServer GS140 Systems: Booting Problem Fixed	6-9
6.6.7	AlphaServer GS60/GS60E/GS140 Multiple I/O Port Module Configuration Restriction	6-9
6.7	AlphaStation 200/400: ISA_CONFIG.DAT Changes Required	6-9
6.8	AlphaStation 255: PCI Configuration Restriction	6-10
6.9	ATI RADEON 7500 Graphics	6-10
6.9.1	Resource Requirements	6-10
6.9.2	Support Limitations	6-11
6.9.3	Message When AGP Card Not Found	6-11
6.9.4	Cloned Video Restricted to Primary Head	6-11
6.9.5	Monitor Must Be Connected During Initialization	6-12
6.9.6	Boot Reset Recommendation	6-12
6.9.7	No Overlay Planes	6-12
6.9.8	Single Colormap	6-12
6.9.9	Single Bit Depth for All Windows	6-12
6.9.10	Pixel Depth for Read/Write Color Map	6-12
6.9.11	Backing Store/Save Unders Not Supported for 3D Windows	6-13

6.9.12	Threads Restriction	6-13
6.9.13	No Support for Single Buffered Visuals	6-13
6.9.14	No 3D Support for Color Index Mode	6-13
6.9.15	Timer Mechanism	6-13
6.10	DECwindows X11 Display Server	6-14
6.10.1	S3 Multihead Graphics	6-14
6.10.2	Integrated Graphics Board Support	6-14
6.11	DIGITAL Modular Computing Components (DMCC)	6-14
6.11.1	Alpha 5/366 and 5/433 PICMG SBC Restriction	6-14
6.11.2	Updating the SRM Console	6-15
6.12	Digital Personal Workstation: Booting OpenVMS V7.3-1 and Higher	6-15
6.13	Dual-Controller HSGnn with Many LUNs Can Fail Under Heavy I/O Load	6-15
6.14	PowerStorm 300/350 PCI Graphics Support for OpenVMS	6-16
6.15	RFnn DSSI Disk Devices and Controller Memory Errors	6-16
6.16	RZnn Disk Drive Considerations	6-18
6.16.1	RZ25M and RZ26N Disk Drives: Recommendations	6-18
6.16.2	RZ26N and RZ28M Disks: Recommended Firmware Support	6-19
6.16.3	RZ26L and RZ28 Disks: Required Firmware for Multihost Use	6-19
6.16.3.1	Firmware Revision Level 442 Requirements	6-19
6.16.3.2	Firmware Revision Level 442 Installation Procedure	6-19
6.17	ZLX Graphics Boards Support	6-20
6.18	Recompiling and Relinking OpenVMS Device Drivers	6-20
6.18.1	Alpha and VAX SCSI Device Drivers	6-20
6.18.2	OpenVMS Alpha Device Drivers	6-20
6.19	Device Driver MON Version Handling	6-21
6.20	Possible Per-Threads Security Impact on Alpha Device Drivers	6-21
6.21	Device IPL Setup for OpenVMS Alpha Drivers	6-21
6.22	CRCTX Routines Enhanced	6-21

A Product Retirement Notices

A.1	Attunity Connect "On Platform" Package	A-1
A.2	ISA_CONFIG.DAT Unsupported in Future Release	A-1
A.3	POSIX 1003.4a Draft 4 Interface May Be Retired	A-2
A.4	Archived Manuals	A-2
A.4.1	Extended File Specifications: Documentation Reorganization	A-2

B Interlocked Memory Instructions

B.1	Required Code Checks	B-1
B.2	Using the Code Analysis Tool (SRM_CHECK)	B-1
B.3	Noncompliant Code Characteristics	B-2
B.4	Coding Requirements	B-3
B.5	Compiler Versions	B-5
B.6	Recompiling Code with ALONONPAGED_INLINE or LAL_REMOVE_FIRST	B-6

Index

Tables

4-1	Patch Kits Required for Cluster Compatibility.....	4-7
4-2	TFF Character Fallback Tables	4-21
6-1	Changes to Device Description Block	6-10
6-2	Supported Microcode Revision Levels	6-17
6-3	Commands for Updating Microcode in Certain DSSI Disk Devices ...	6-18
6-4	Revision Level 442 Firmware Compatibility	6-19
B-1	Versions of OpenVMS Compilers.....	B-5

Preface

Intended Audience

This manual is intended for all users of the HP OpenVMS Alpha Version 7.3-2 operating system. Read this manual before you install, upgrade, or use OpenVMS Alpha Version 7.3-2.

Note

Please check the *HP OpenVMS Alpha Version 7.3-2 Cover Letter* for release notes and late-breaking information that was not available at the time this manual went to production. You can access the cover letter on the OpenVMS documentation web site:

<http://www.hp.com/go/openvms/doc>

Document Structure

This manual contains the following chapters and appendixes:

- Chapter 1 contains release notes that pertain to installing or upgrading the OpenVMS Alpha operating system.
- Chapter 2 contains installation and support information for OpenVMS Alpha associated products.
- Chapter 3 contains release notes about the general use of the OpenVMS Alpha operating system.
- Chapter 4 contains release notes specific to OpenVMS system management.
- Chapter 5 contains release notes that relate to programming on an OpenVMS Alpha system, including notes for compilers, linkers, and run-time library routines.
- Chapter 6 contains information pertaining to hardware that runs on the OpenVMS Alpha operating system and OpenVMS device support on Alpha systems.
- Appendix A contains information about OpenVMS Alpha products that are no longer supported as of this release, or that are slated for retirement.
- Appendix B describes the proper use of interlocked memory instructions, which is crucial for the Alpha 21264 (EV6) processor.

Notes are organized by facility or product name when applicable.

This manual contains release notes introduced in the current release and notes from previous OpenVMS Alpha versions that still apply to the new release. A subheading for each release note indicates either the version of origin (for example, V7.3) or the version when an old note was last updated (for example, a note from V7.3 that was revised for Version 7.3-2 will be labeled with V7.3-2).

Notes from previous releases are published when:

- The information in the release note has not been documented in any other printed manual in the OpenVMS documentation set, and the note is still pertinent.
- The release note may be pertinent in multiple-version OpenVMS Cluster systems.

Related Documents

For a list of additional documents that are available in support of this version of the OpenVMS operating system, refer to the *HP OpenVMS Alpha Version 7.3-2 New Features and Documentation Overview* manual.

For additional information about HP OpenVMS products and services, visit the following World Wide Web address:

<http://www.hp.com/go/openvms>

Reader's Comments

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

Internet	openvmsdoc@hp.com
Postal Mail	Hewlett-Packard Company OSSG Documentation Group, ZKO3-4/U08 110 Spit Brook Rd. Nashua, NH 03062-2698

How to Order Additional Documentation

For information about how to order additional documentation, visit the following World Wide Web address:

<http://www.hp.com/go/openvms/doc/order>

Conventions

The following conventions may be used in this manual:

Ctrl/x	A sequence such as Ctrl/x indicates that you must hold down the key labeled Ctrl while you press another key or a pointing device button.
PF1 x	A sequence such as PF1 x indicates that you must first press and release the key labeled PF1 and then press and release another key or a pointing device button.

Return

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.

...

A horizontal ellipsis in examples indicates one of the following possibilities:

- 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.

.
. .
.

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.

()

In command format descriptions, parentheses indicate that you must enclose choices in parentheses if you specify more than one.

[]

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.

|

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.

{ }

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 type

Bold type represents the introduction of a new term. It also represents the name of an argument, an attribute, or a reason.

italic type

Italic type indicates important information, complete titles of manuals, or variables. Variables include information that varies in system output (Internal error *number*), in command lines (/PRODUCER=*name*), and in command parameters in text (where *dd* represents the predefined code for the device type).

UPPERCASE TYPE

Uppercase type indicates a command, the name of a routine, the name of a file, or the abbreviation for a system privilege.

Example

This typeface indicates code examples, command examples, and interactive screen displays. In text, this type also identifies URLs, UNIX commands and pathnames, PC-based commands and folders, and certain elements of the C programming language.

-

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.

OpenVMS Alpha Software Installation and Upgrade Release Notes

This chapter contains information that you need to know before installing or upgrading to the OpenVMS Alpha Version 7.3-2 operating system.

HP recommends that you read all of the following manuals before installing or upgrading to OpenVMS Alpha Version 7.3-2:

- *HP OpenVMS Alpha Version 7.3-2 Release Notes* (this manual)
- *HP OpenVMS Alpha Version 7.3-2 New Features and Documentation Overview*
- *HP OpenVMS Alpha Version 7.3-2 Upgrade and Installation Manual*

Also refer to Chapter 6 for hardware release notes and to Chapter 2 for information about associated products.

1.1 HP Software Technical Support Policy

V7.3-2

HP provides software technical support for OpenVMS operating system software for the latest, currently shipping version and the immediate prior version of the product. Each version is supported for 24 months from its release date, or until the release of the second subsequent version, whichever is greater. “Version” is defined as a release containing new features and enhancements. Patch kits and maintenance-only releases do not meet the definition of “version” in the context of this support policy.

Current version-level support (Standard Support or SS) and Prior Version Support (PVS) for OpenVMS operating system software is provided for OpenVMS versions in accordance with these guidelines. The current level of support for recent versions of OpenVMS Alpha and OpenVMS VAX is kept up to date on line. You can access current support information by following the links on this web page:

http://h18005.www1.hp.com/services/software/ss_openvms.html

The Operating System Support Policy applies to all OpenVMS Major Releases, New Feature Releases, and Enhancement releases, which are defined as follows:

- **Major Releases** for OpenVMS contain substantial new functionality. The version number increases to the next integer (for example, from 6.2-1H1 to 7.0).

Application impact: OpenVMS internal interfaces have changed. Although binary compatibility will be maintained for the majority of applications, independent software vendors (ISVs) should test on the new version and may need to release a new application kit. Some application partners may want

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.1 HP Software Technical Support Policy

to release a new application kit to take advantage of new features in the operating system.

- **New Feature Releases** for OpenVMS contain new features, enhancements, and maintenance updates. The version number increases to the next decimal number (for example, from 7.2 to 7.3).

Application impact: The release has not retired any published application programming interfaces (APIs). However, OpenVMS internal interfaces may have been modified with the addition of significant new functionality or implementation of performance improvements. It is unlikely that a new application kit will be required for 95 percent of all applications that use documented APIs. Device driver and kernel-level applications (that is, those that use nonstandard or undocumented APIs) may need qualification testing.

- **Enhancement Releases** for OpenVMS contain enhancements to existing features and maintenance updates. The version number increases to show a revision by using a dash (for example, OpenVMS Version 7.3-1).

Application impact: The release may contain new hardware support, software enhancements, and maintenance, but the changes are isolated and have no impact on applications that use published APIs. There is no need for ISVs to test on the new release or to produce a new application kit.

- **Hardware Releases** provide current version-level support until 12 months after a subsequent release containing that particular hardware support. Hardware releases are shipped with new hardware sales only and are not distributed to existing service customers.

The following OpenVMS core products are supported at the same level (Standard Support or Prior Version Support) and duration as the operating system version on which they shipped:

- HP Advanced Server for OpenVMS
- HP DECnet (Phase IV)
- HP DECnet-Plus for OpenVMS
- HP OpenVMS Cluster Client Software
- HP OpenVMS Cluster Software for OpenVMS
- HP PathWorks or HP PATHWORKS for OpenVMS
- HP RMS Journaling for OpenVMS
- HP TCP/IP Services for OpenVMS
- HP Volume Shadowing for OpenVMS
- HP X.25 for OpenVMS

These products require their own individual support contracts and are not included in the operating system support contract.

1.2 General Application Compatibility Statement

OpenVMS has consistently held the policy that published APIs are supported on all subsequent releases. It is unlikely applications that use published APIs will require changes to support a new release of OpenVMS. APIs may be "retired," and thus removed from the documentation; however, the API will continue to be available on OpenVMS as an undocumented interface.

1.3 Firmware for OpenVMS Alpha Version 7.3-2

V7.3-2

OpenVMS Alpha Version 7.3-2 was tested with the platform-specific firmware included on Alpha Systems Firmware CD-ROM Version 6.5. For older platforms no longer included on the Firmware CD-ROM, OpenVMS Version 7.3-2 was tested with the latest released firmware version. HP recommends upgrading to the latest firmware before upgrading OpenVMS.

The OpenVMS Alpha Version 7.3-2 CD-ROM package includes the Alpha Systems Firmware CD-ROM and Release Notes. Read the Release Notes before installing the firmware.

For AlphaServer ES47/ES80/GS1280 systems, use Alpha Systems Firmware Version 6.6 or later. You can obtain Version 6.6 and the latest firmware information from the following website (URL is case sensitive):

<ftp://ftp.digital.com/pub/Digital/Alpha/firmware/>

1.4 Upgrade Paths

V7.3-2

Upgrade Paths to Version 7.3-2

You can upgrade *directly* to OpenVMS Alpha Version 7.3-2 from only the following versions of OpenVMS Alpha:

- Version 7.3-1
- Version 7.3
- Version 7.2-2
- Version 7.2-1
- Version 7.2-1H1
- Version 7.2

If you are currently running OpenVMS Alpha Version 6.2x through 7.1x, you must do a two-step upgrade: first to Version 7.2-2 or 7.3, and then to Version 7.3-2. If you are running other versions of OpenVMS that are no longer supported, you will have to do multiple upgrades in accordance with upgrade paths that were documented for earlier versions.

Cluster Concurrent Upgrades

During a concurrent upgrade, you must shut down the entire cluster and upgrade each system disk. No one can use the cluster until you upgrade and reboot each computer. Once you reboot, each computer will be running the upgraded version of the operating system.

Cluster Rolling Upgrades

During a cluster rolling upgrade, you upgrade each system disk individually, allowing old and new versions of the operating system to run together in the same cluster (a mixed-version cluster). There must be more than one system disk. The systems that are not being upgraded remain available.

Only the following OpenVMS Alpha and OpenVMS VAX versions are supported in mixed-version clusters that include OpenVMS Alpha Version 7.3-2:

- Version 7.3-1 (Alpha)
- Version 7.3 (Alpha and VAX)
- Version 7.2-2 (Alpha)

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.4 Upgrade Paths

Version 7.2 (VAX)

If you are upgrading in a cluster environment, rolling upgrades are supported from Version 7.2-2, 7.3, and 7.3-1 of the OpenVMS Alpha operating system. If you have other versions in a cluster, you cannot do a rolling upgrade until those versions are upgraded to a supported version.

Mixed-version support for some of these versions is restricted for cluster members running either the Registry server or applications that use the Registry services. See Section 1.9.10.

Mixed-version support for all of these versions requires the installation of one or more remedial kits. For more information, see Section 4.13.1.

Note

HP currently supports only two versions of OpenVMS (regardless of architecture) running in a cluster at the same time.

For a discussion of warranted pairs and migration pairs of OpenVMS operating systems, as well as complete instructions for installing or upgrading to OpenVMS Alpha Version 7.3-2, refer to the *HP OpenVMS Alpha Version 7.3-2 Upgrade and Installation Manual*.

1.4.1 Cluster Rolling Upgrades: Correction to *HP OpenVMS Alpha Version 7.3-2 Upgrade and Installation Manual*

V7.3-2

In Chapter 5 of the *HP OpenVMS Alpha Version 7.3-2 Upgrade and Installation Manual*, there is an omission in the Cluster Rolling Upgrades section. The list of versions supported with OpenVMS Alpha Version 7.3-2 in mixed-version clusters is incomplete. The list should also include OpenVMS Version 7.2 (VAX).

1.5 Networking Options

V7.3-2

OpenVMS provides customers with the flexibility to choose the network protocol of their choice. Whether you require DECnet or TCP/IP, OpenVMS allows you to choose the protocol or combination of protocols that works best for your network. OpenVMS can operate with both HP and third-party networking products.

During the main installation procedure for OpenVMS Version 7.3-2, you have the option of installing the following supported HP networking software:

- Either HP DECnet-Plus Version 7.3-2 for OpenVMS or HP DECnet Phase IV Version 7.3-2 for OpenVMS. (Note that both DECnet products cannot run concurrently on your system.)

DECnet-Plus contains all the functionality of the DECnet Phase IV product, plus the ability to run DECnet over TCP/IP or OSI protocols.

Support for DECnet Phase IV is provided to customers with a Prior Version Support Contract. For more information about the Prior Version Support service, see Section 1.1.

- HP TCP/IP Services for OpenVMS Version 5.4

TCP/IP Services and DECnet can run concurrently on your system. Once you have installed HP DECnet-Plus for OpenVMS and TCP/IP Services on your system, you can run DECnet applications and OSI applications, or both, over your TCP/IP network. Refer to the *DECnet-Plus for OpenVMS Management Guide* for more information about running DECnet over TCP/IP (RFC 1859) and OSI over TCP/IP (RFC 1006).

Alternatively, after you install OpenVMS, you can install your choice of another third-party networking product that runs on OpenVMS Version 7.3-2.

For information about how to configure and manage your HP networking software after installation, refer to the TCP/IP, DECnet-Plus, or DECnet documentation. The manuals are available in online format on the OpenVMS Documentation CD-ROM; printed manuals can be ordered from HP (800-282-6672).

1.6 Obtaining Remedial Kits

V7.3-2

Remedial kits for HP products are available online at the HP IT Resource Center (ITRC). Use of the ITRC patch download site requires user registration and login. Registration is open to all users and no service contract is required. You can register and log in from the following URL:

<http://www2.itrc.hp.com/service/patch/mainPage.do>

You can also use FTP to access patches from the following location:

ftp://ftp.itrc.hp.com/openvms_patches

1.7 System Startup Failure in a Cluster

V7.3-2

During system startup, SET VOLUME/SUBSYSTEM is performed on the system disk. This operation might silently fail in a cluster environment if the system disk is locked (for example, during a rebuild operation). If SET VOLUME/SUBSYSTEM fails, a second attempt is made to enable the subsystem after the call to SYSTARTUP_VMS.COM. If the lock still has not been released and the operation fails again, the following warning is issued:

```
%STARTUP-W-SUBSYSNOTENA, failed to enable subsystem ACEs on system disk
```

1.8 SHADOW_MAX_UNIT Default Setting and Memory Usage

V7.3-2

This note updates an earlier note that discussed the default settings for this system parameter but did not describe the amount of main memory consumed by the default settings.

OpenVMS Alpha Version 7.3 introduced minicopy support in HP Volume Shadowing for OpenVMS. As part of the minicopy functionality, a new volume shadowing system parameter, SHADOW_MAX_UNIT, was introduced. On OpenVMS Alpha systems, the default value for this system parameter is 500, which consumes 24 KB of main memory. On OpenVMS VAX systems, the default value is 100, which consumes 5 KB of main memory.

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.8 SHADOW_MAX_UNIT Default Setting and Memory Usage

If you do not plan to use Volume Shadowing for OpenVMS, you can change the setting to its minimum of 10 (which consumes 480 bytes of main memory). By setting the default to its minimum, you free up 23.5 KB of main memory on an OpenVMS Alpha system and 4.5 KB of main memory on a VAX system.

Note

SHAD_MAX_UNIT is a static system parameter. In order for a new setting to take effect, you must reboot your system.

Recommendations for SHADOW_MAX_UNIT settings for volume shadowing are discussed in the *HP Volume Shadowing for OpenVMS* manual.

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

The remainder of this chapter concerns associated products that might affect the installation or upgrade of the OpenVMS operating system. See Chapter 2 for more associated products release notes.

1.9.1 Some Product Kits Ship in Compressed Format

V7.3-2

Starting with OpenVMS Version 7.3-2, some of the POLYCENTER Software Installation utility kits on the operating system and layered product CD-ROMs are provided in compressed format. You can identify these kits by their file types:

- Compressed sequential kits have a .PCSI\$COMPRESSED extension.
- Uncompressed sequential kits have a .PCSI extension.

The compressed format was developed to save space on the CD-ROMs. A bonus is that user systems with a slow CD-ROM reader may experience a slightly faster installation time for compressed kits because less data must be transferred from the input device.

The installation of a compressed kit is completely transparent to the user; no qualifiers or special actions are required. The POLYCENTER Software Installation utility expands (decompresses) records on-the-fly as it installs the product, so no extra working disk space is required.

The POLYCENTER Software Installation utility supplied with OpenVMS Version 7.3-2 (and any later versions) can use compressed and uncompressed kits interchangeably. However, on versions of OpenVMS prior to 7.3-2, the compressed (.PCSI\$COMPRESSED) kits cannot be read unless an appropriate POLYCENTER Software Installation utility remedial kit is installed. This remedial kit is expected to be available before the final release of OpenVMS Version 7.3-2.

Converting a Compressed Kit to an Uncompressed Kit

As an alternative to installing a POLYCENTER Software Installation utility remedial kit on an earlier system, you can easily convert a compressed kit into an uncompressed format. You can use the PRODUCT COPY command on OpenVMS Version 7.3-2 to create a standard .PCSI kit from a .PCSI\$COMPRESSED kit and then transfer it to the target system for installation. Use a command similar to the following to perform the conversion:

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

```
$ PRODUCT COPY product-name /SOURCE=device:[dir] -  
_$_ /DESTINATION=device:[dir] /FORMAT=SEQUENTIAL
```

Not all kits provided on the OpenVMS Version 7.3-2 operating system and layered products CD-ROMs can be installed on prior versions of OpenVMS. Before you attempt to install a product on a prior version, consult the product's documentation to determine whether this kit installation is supported.

1.9.2 ACMS Kits: Problem with File Deletions

V7.3-1

If you have installed any of the following kits:

```
VMS73_ACMS-V0100  
VMS722_ACMS-V0100  
VMS721H1_ACMS-V0100  
VMS721_ACMS-V0100
```

and then you upgrade to Version 7.3-1 or higher, the following ACMS files might get deleted:

```
ACMSTART.COM  
ACMSBOOT.EXE
```

This will cause ACMS to fail after the upgrade. To avoid this problem (or recover from it), remove the old ACMS kit (select Option 6 "Remove installed products" on the installation CD main menu) and then install the ACMS_U1_044 kit either before or after you upgrade the operating system. This kit replaces the kits in the preceding list.

You can get the latest ACMS kit from the following web site (follow the version links to the ACMS_U1_044 kit):

<http://ftp1.support.compaq.com/public/vms/axp/>

1.9.3 COM for OpenVMS Upgrade in Clusters Running OpenVMS Version 7.2-1 or 7.2-1H1

V7.3-1

If you are running COM for OpenVMS in a cluster running OpenVMS Alpha Version 7.2-1 or Version 7.2-1H1, and you intend to do a rolling upgrade, see Section 1.9.10 for information about how to upgrade your cluster.

1.9.4 DECevent Version 3.1 or Higher Required for Certain Systems

V7.3-1

Analyzing Earlier Hardware Platforms

DECevent Version 3.1 or higher is required to analyze hardware error log files on certain supported, earlier hardware platforms (for example, the 1200, 8400, GS60, and GS140 platforms).

When you install or upgrade to OpenVMS Version 7.3 or higher, the DECevent DCL command DIAGNOSE is disabled. If you are installing OpenVMS on a hardware platform that needs DECevent and its DIAGNOSE command, perform the following steps:

1. Install or upgrade to OpenVMS Version 7.3 or higher.
2. Install the DECevent software (included in the DECevent kit on the HP System Tools CD-ROM).

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

Otherwise, when you attempt to use the DIAGNOSE command, you will receive the following system message:

```
$ DIAGNOSE [parameters]
%DIA-E-NOINSTAL, DIAGNOSE has not been installed on this system
```

Analyzing Later Platforms

The System Event Analyzer (SEA) is now the supported error log analysis tool for OpenVMS and for later hardware platforms (for example, the DSnn, ESnn, GS80, GS160, and GS320 platforms).

Additionally, you can use the new Error Log Viewer (ELV) utility to quickly examine an error log file that was created on these later hardware platforms. ELV is integrated into OpenVMS Version 7.3-2, and can be accessed by entering the DCL command ANALYZE/ERROR_LOG/ELV.

For more information about ELV, refer to the ELV chapter in *HP OpenVMS System Management Utilities Reference Manual*.

Analyzing Both Platforms

If you have both a DECEvent-supported platform and SEA-supported storage devices, you will need both DECEvent and the System Event Analyzer (SEA). Likewise, if you have both an SEA-supported platform and DECEvent-supported storage or other devices, you will need both tools.

To install SEA, you must install the latest version of WEBES (Version 4.2 or higher). The current version of the DECEvent kit is Version 3.4.

For detailed information about operating system requirements and supported hardware for DECEvent, refer to the DECEvent Release Notes, which are at the following web site:

<http://h18023.www1.hp.com/support/svctools/decevent/>

Follow the "Download the documentation" link.

For detailed information about operating system requirements and supported hardware for SEA, refer to the WEBES Installation Guide, which is with the other WEBES documentation at the following web site:

<http://h18023.www1.hp.com/support/svctools/webes/>

For more information about DECEvent, refer to the *HP OpenVMS System Manager's Manual, Volume 2: Tuning, Monitoring, and Complex Systems*.

1.9.5 DECnet-Plus and DECwindows Require New Versions

V7.3-2

When you install or upgrade to OpenVMS Alpha Version 7.3-2, you must also install new versions of DECwindows and DECnet-Plus. One of the reasons that makes this necessary is a change of behavior in AUTOGEN (see Section 4.2).

Unlike the behavior of previous versions, DECnet-Plus for OpenVMS Version 7.3-2 now provides product information in NEWPARAMS.DAT records, as required by AUTOGEN. AUTOGEN anticipates this change in DECnet-Plus, so AUTOGEN does not print any warnings when it removes "bad" records from CLU\$PARAMS.DAT; AUTOGEN presumes these records were made by an older DECnet-Plus kit and will be replaced by the new DECnet-Plus kit. So, under normal conditions, you will not see any striking differences in behavior during an OpenVMS Version 7.3-2 installation or upgrade.

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

However, if other products do not provide product information in NEWPARAMS.DAT records, as now required by AUTOGEN, AUTOGEN prints warning messages to both the report and the user's SYS\$OUTPUT device. The warnings state that AUTOGEN cannot accept the parameter assignment found in NEWPARAMS.DAT (because no product name is attached) and that no records will be added to CLU\$PARAMS.DAT. Since no records are added, the expected additions or other alterations to SYSGEN parameters will not be made, which could lead to resource exhaustion. Developers and testers of software products should be aware of this requirement; it may also be of interest to system managers.

This new behavior is intended to protect both the users and providers of layered products. By keeping this information ordered properly so that it can be updated properly, problems resulting from bad updates should be minimized.

A description of NEWPARAMS.DAT and CLU\$PARAMS.DAT is included in the AUTOGEN chapter of the *HP OpenVMS System Management Utilities Reference Manual*.

1.9.6 DECnet-Plus Installation: PCSI-I-RETAIN Messages

V7.2

If you upgrade to OpenVMS Version 7.3 or higher *and* your system has either DCE for OpenVMS or DECnet-Plus for OpenVMS installed on it, when you install DECnet-Plus you may get PCSI-I-RETAIN informational messages for the following files:

```
[SYSEXEC]DTSS$SET_TIMEZONE.EXE
[SYSLIB]DTSS$RUNDOWN.EXE
[SYSUPD]DTSS$TIMEZONE_RULES.DAT
[SYSLIB]DTSS$SHR.EXE
```

For example:

```
%PCSI-I-RETAIN, file [SYSEXEC]DTSS$SET_TIMEZONE.EXE was not replaced
because file from kit has a lower generation number
```

You can ignore these messages. The DECnet-Plus kit has been properly installed.

1.9.7 HP PATHWORKS and HP Advanced Server for OpenVMS Products

These notes pertain to HP PATHWORKS and HP Advanced Server for OpenVMS products, and include information about installing or upgrading OpenVMS systems running these products.

1.9.7.1 HP Advanced Server for OpenVMS

V7.3-2

Version 7.3A of Advanced Server for OpenVMS is supported on OpenVMS Alpha Version 7.3-1 and higher systems. Advanced Server Versions 7.2 and 7.2A for OpenVMS servers must be upgraded. For more information about upgrading Advanced Server for OpenVMS servers, see Section 1.9.7.5.

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

1.9.7.2 HP PATHWORKS for OpenVMS (Advanced Server)

V7.3-1

Version 6.1 of PATHWORKS for OpenVMS (Advanced Server) is supported on OpenVMS Alpha Version 7.3-1 and higher systems. Earlier versions of PATHWORKS for OpenVMS servers must be upgraded. For more information about upgrading earlier versions of PATHWORKS, see Section 1.9.7.4.

1.9.7.3 PATHWORKS V5 for OpenVMS (LAN Manager) Not Supported

V7.3-2

PATHWORKS V5 for OpenVMS (LAN Manager) is not supported on OpenVMS Version 7.3-2.

If you are running PATHWORKS V5 for OpenVMS (LAN Manager) and you want to offer file and print services on OpenVMS Version 7.3-2, you must upgrade the file and print server to PATHWORKS V6.1 for OpenVMS (Advanced Server) before you install OpenVMS Version 7.3-2.

You cannot upgrade directly from PATHWORKS V5 for OpenVMS (LAN Manager) to Advanced Server V7.3x for OpenVMS. For information about upgrading from PATHWORKS V5 for OpenVMS (LAN Manager) to PATHWORKS V6.x for OpenVMS (Advanced Server), refer to the *PATHWORKS for OpenVMS (Advanced Server) Server Installation and Configuration Guide* provided with the kit documentation. For information about upgrading earlier versions of PATHWORKS for OpenVMS (Advanced Server) to PATHWORKS V6.1 for OpenVMS (Advanced Server), see Section 1.9.7.4.

1.9.7.4 Upgrading Systems with Pre-V6.1 PATHWORKS Advanced Server

V7.3-2

If you are upgrading an OpenVMS system that is currently running a version of PATHWORKS for OpenVMS (Advanced Server) that is earlier than Version 6.1, follow these steps:

1. Upgrade your PATHWORKS for OpenVMS (Advanced Server) to Version 6.1.
2. Upgrade your OpenVMS system to OpenVMS Version 7.3-2.
3. At this point, you also have the option of upgrading PATHWORKS for OpenVMS (Advanced Server) to Advanced Server V7.3A for OpenVMS.

1.9.7.5 Upgrading Advanced Server Version 7.2x for OpenVMS

V7.3-2

If you want to upgrade your Advanced Server for OpenVMS server, follow these steps:

1. Upgrade your Advanced Server V7.2 or V7.2A for OpenVMS server to Advanced Server V7.3A for OpenVMS.
2. Upgrade your OpenVMS Alpha system to OpenVMS Version 7.3-2.

Note

Because of changes to the OpenVMS Registry protocol, you cannot run Advanced Server for OpenVMS software on OpenVMS Alpha Version 7.2-2 (and higher) systems *and* on OpenVMS Alpha systems prior to Version 7.2-2 in the same cluster.

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

For more information about upgrading systems in mixed-version clusters, see Section 1.9.10.

1.9.8 HP X.25 Version 1.6 Required Upgrade

V7.3-1

Customers with HP X.25 for OpenVMS Alpha systems software must upgrade to X.25 Version 1.6 prior to upgrading to OpenVMS 7.3-1 or higher. Failure to upgrade will result in system failures with an SPLINVIPL bugcheck when booting.

HP also recommends that you install the latest X.25 ECO kit, which is named X25ALP X25_V16ECO1 X.25 V1.6 for OpenVMS Alpha. This ECO kit contains two separate kits, one for X.25 and one for WANDD, which are named as follows:

- DEC-AXPVMS-X25-V0106-1-1.PCSI
- DEC-AXPVMS-WANDD-V0106-1-1.PCSI

You can access the X25_V16ECO1 kit from the following website if you are a contract customer and already have an account:

<http://ftp.support.compaq.com/cgi-bin/entitlement.cgi/patches/entitled/>

If you are not already a contract customer or if you do not have an account at the patch site, you can create an account at this site or contact a customer service representative to obtain the kit.

1.9.9 Kerberos V1.0: Remove Before Upgrading

V7.3-2

If you installed Kerberos Version 1.0 for OpenVMS using a POLYCENTER Software Installation kit, you must use the POLYCENTER Software Installation utility to remove Kerberos Version 1.0 before you upgrade the operating system. (You do *not* need to remove Kerberos if you are running Version 2.0 or if Version 1.0 was installed as part of the OpenVMS Version 7.3-1 operating system.)

To remove Kerberos, choose Option 6 “Remove installed products” from the installation CD main menu. During the removal, you are asked whether you want to remove the data and directories. (Data refers to the configuration data files along with the principal database, if one was created.) If you want to save this information for use later, respond “No” to the question. Return to the main menu and perform the upgrade of OpenVMS.

After the upgrade, the new Kerberos directories are located in KRB\$ROOT:[*...]. (KRB\$ROOT is defined as a system-wide logical name when Kerberos is started.) Kerberos data is either created during configuration or moved from the old Kerberos directories and renamed. If you removed a previously installed Kerberos kit and saved the data and directories, the data will automatically be moved into the new directories and be renamed the first time the Kerberos startup procedure is run after the upgrade.

Start the Kerberos servers by entering the following command:

```
$ @sys$startup:krb$startup.com
```

Note that the Kerberos startup procedure moves and renames only known Kerberos files. Users who have created files in the old Kerberos directories must manually move those files.

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

For more information about installing and configuring Kerberos, refer to *HP Open Source Security for OpenVMS, Volume 3: Kerberos*.

1.9.10 Registry Upgrade from OpenVMS Version 7.2-1 or 7.2-1H1

V7.3-1

The Registry that shipped with OpenVMS Alpha Version 7.2-1 and Version 7.2-1H1 is different from and incompatible with the Registry that ships with Versions 7.2-2 and higher. If you are upgrading from an earlier version with the old Registry, you must take special steps, depending on what you want to do.

- **If you choose to upgrade all Alpha nodes in the cluster at once:**
Shut down only the Registry and all applications using the Registry before upgrading, and then reverse the process for startup after upgrading.
- **If you choose to upgrade only some nodes in the cluster at a time:**
You can run Registry servers and applications on only the Version 7.2-1 and 7.2-1H1 nodes in the cluster — or on only the Version 7.2-2, 7.3, 7.3-1, and 7.3-2 nodes in the cluster — but not on both. Before you upgrade each node in the cluster, you must inhibit the startup of the following on the upgraded node:
 - The Registry
 - Advanced Server
 - COM for OpenVMS
 - Any other applications using the Registry

At some point, you must shut down all remaining Registry-based activity in the cluster just before you start up Registry and applications using Registry services on the Version 7.2-2 and higher nodes.

Note

If you are running Advanced Server V7.2 or V7.2A for OpenVMS on OpenVMS Version 7.2-1 or 7.2-1H1, you must upgrade all nodes to Advanced Server V7.3 for OpenVMS before you upgrade any OpenVMS Version 7.2-1 or 7.2-1H1 node to OpenVMS Version 7.3 or higher. If you are upgrading to OpenVMS Version 7.3-2, once you complete the upgrade, you must upgrade Advanced Server for OpenVMS again to V7.3A or higher.

The following steps describe the procedure you can use when upgrading from Version 7.2-1 or from 7.2-1H1 to Version 7.2-2 or higher on systems running the OpenVMS NT Registry:

1. Though not required, it is best to shut down the Registry in a graceful manner. Before shutting down the Registry, shut down all layered products that use the Registry. First, shut down any applications specific to your environment that are known to use Registry services. Then shut down HP layered products that use Registry services. For example:

- First shut down COM for OpenVMS using the following command:

```
$ @SYS$STARTUP:DCOM$SHUTDOWN.COM
```

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

- Then shut down Advanced Server using this command:

```
$ @SYS$STARTUP:PWRK$SHUTDOWN.COM
```

2. Create a snapshot of the Registry database by using the following commands:

```
$ REG$CP ::= $REG$CP  
$ REG$CP CREATE SNAPSHOT
```

3. Export the Registry database by using the command:

```
$ REG$CP EXPORT DATABASE [/LOG/OUTPUT=filename]
```

4. If you are upgrading all nodes in the cluster at the same time, make a note as to which node is acting as the master Registry server. You can determine which node is the master by issuing the command:

```
$ SHOW SERVER REGISTRY/MASTER
```

5. Shut down the Registry server or servers. If you are upgrading all nodes in the cluster at the same time, this can be performed using the command:

```
$ SET SERVER REGISTRY/CLUSTER/EXIT
```

If you are upgrading just one node in the cluster, issue the following command on the node:

```
$ SET SERVER REGISTRY/EXIT
```

If that node is the master, wait until it exits before you take any other action. Another node in the cluster will become the master.

6. Ensure that the Registry server does not restart on the node or nodes you are upgrading until the upgrade is complete, or, if you are selectively upgrading nodes, until you determine that you wish to switch over to the new server.

To prevent Registry startup on reboot, you need to check two things on each node:

- a. In the file `SYS$MANAGER:SYLOGICALS.COM`, either comment out or set to FALSE any logical name definitions that contain the string:

```
"TO_BE_STARTED"
```

Make a note of the original settings for restoring later.

- b. If your `SYS$MANAGER:SYSTARTUP_VMS.COM` file includes commands that automatically start up layered products that use Registry services, comment those lines out of the file so that those products are not automatically started on that node. Look for lines such as the following, which starts up the Advanced Server:

```
$ @SYS$STARTUP:PWRK$STARTUP.COM
```

7. Proceed with the upgrade on each node. (For detailed upgrade information, refer to the *HP OpenVMS Alpha Version 7.3-2 Upgrade and Installation Manual*.)

8. Once all nodes have been upgraded, restart the master server by using the following command on the node that was originally running the master server:

```
$ SET SERVER REGISTRY/START
```

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

If you are selectively upgrading nodes, and you are ready to switch to using Registry services on the upgraded nodes, shut down the Registry server, and applications using Registry services, on all remaining OpenVMS Version 7.2-1 and 7.2-1H1 nodes in the cluster using steps 1-6 outlined above. Then you can start the Registry server on one of the upgraded nodes.

9. Verify that the Registry is operational by using the following commands:

```
$ REG$CP :== $REG$CP
$ REG$CP LIST KEY HKEY_LOCAL_MACHINE
```

The last command should display at least four subkeys of the HKEY_LOCAL_MACHINE root key. The same command should be repeated with the HKEY_USERS root key, which should display at least one subkey.

Note

In the unlikely event that the Registry is not operational, follow the steps in the *HP COM, Registry, and Events for OpenVMS Developer's Guide* describing how to restore your database from the snapshot files. If this fails, delete all the files in the SYS\$REGISTRY directory, or rename the directory, and invoke SYS\$STARTUP:REG\$CONFIG to reconfigure the Registry server (refer to the *HP COM, Registry, and Events for OpenVMS Developer's Guide* for details), then import the database file that was saved in step 3.

10. Start the backup Registry servers on the other upgraded nodes using the command:

```
$ SET SERVER REGISTRY/START
```

11. Restore the values of "TO_BE_STARTED" logical name definitions in SYS\$MANAGER:SYLOGICALS.COM and the invocation of Advanced Server startup in the SYS\$MANAGER:SYSTARTUP_VMS.COM file.

If you are selectively upgrading nodes, comment out or set to FALSE the appropriate "TO_BE_STARTED" logical name definitions in the SYS\$MANAGER:SYLOGICALS.COM file and comment out or remove any invocation of Advanced Server startup in the SYS\$MANAGER:SYSTARTUP_VMS.COM file on any remaining OpenVMS Version 7.2-1 nodes in the cluster, as described in step 6.

12. In this order, restart the Advanced Server, COM for OpenVMS, and any other applications that use Registry on the upgraded nodes.

1.9.11 SSL: Installing HP SSL V1.1-A

V7.3-2

If your OpenVMS Alpha system has a version of SSL that is earlier than HP SSL V1.1-A, you must upgrade to HP SSL V1.1-A after you upgrade your operating system to OpenVMS Version 7.3-2 and install TCP/IP Services for OpenVMS. (SSL requires TCP/IP Services, so be sure to select TCP/IP during your system upgrade or installation.) The HP SSL V1.1-A kit ships on the layered products CD-ROM. See the *HP Open Source Security for OpenVMS, Volume 2: HP SSL for OpenVMS* manual for more information.

OpenVMS Alpha Software Installation and Upgrade Release Notes

1.9 Associated Products Affecting OpenVMS Installation and Upgrades

1.9.12 System Event Analyzer (SEA) Utility

V7.3-2

If you run the System Event Analyzer (SEA) for OpenVMS, you must upgrade to SEA Version 4.2 or later to analyze error log files created on systems running OpenVMS Version 7.3-2. Earlier versions of SEA do not support analysis of error log files created on systems running OpenVMS Version 7.3-2.

For detailed information about operating system requirements and supported hardware for SEA, refer to the WEBES Installation Guide, which is posted with the other WEBES documentation on the following web site:

<http://h18023.www1.hp.com/support/svctools/webes/>

OpenVMS Associated Products Release Notes

This chapter contains information about OpenVMS associated products. Notes specifically related to installation or upgrade issues related to associated products are in subsections to Section 2.1.

For notes about using compilers, linkers, and run-time library routines, see Chapter 5.

2.1 Associated Product Support

The Software Public Rollout Reports for OpenVMS list the availability of HP software products shipping on the Software Products Library kits (CD-ROM consolidations) for OpenVMS Alpha and OpenVMS VAX.

The reports contain the product name and version, the operating system version required to support the product, and the volume ship date for the product. The information in these tables is continually evolving and is subject to change. The reports are intended for public distribution and are updated monthly. The information is not provided in these release notes because of the changing nature of the information.

The Software Public Rollout Reports for OpenVMS are available from the following website:

<http://h71000.www7.hp.com/openvms/os/swroll/>

If you do not have Internet access, you can find the operating system support information on any of the quarterly Software Products Libraries in the following files:

[README] SW_COMPAT_MATRIX.PS
[README] SW_COMPAT_MATRIX.TXT

The Software Public Rollout Reports are also available from your HP support representative.

2.2 ACMS: Changes in Error Reporting

V7.3-2

When the Application Control and Management System (ACMS) is run on OpenVMS Alpha Version 7.3-2, there are changes in error reporting under the following conditions:

- The account specified in ACMSGEN for ACC_USERNAME has expired.
Messages reported in prior versions:

OpenVMS Associated Products Release Notes

2.2 ACMS: Changes in Error Reporting

```
$ ACMS/START SYSTEM/NOTERM
%ACMSOPR-E-STRTSYSERR, Error while attempting to START SYSTEM
-ACMSOPS-E-ACCDIED, ACC died unexpectedly.
-LOGIN-F-ACNTEXPIR, your account has expired; contact your system manager
%ACMSOPR-E-ERROR, Some operations may not have been performed
```

Messages reported in OpenVMS Version 7.3-2:

```
$ ACMS/START SYSTEM/NOTERM
%ACMSOPR-E-STRTSYSERR, Error while attempting to START SYSTEM
-ACMSOPS-E-ACCDIED, ACC died unexpectedly.
%ACMSOPR-E-ERROR, Some operations may not have been performed
```

The exact error can be seen in the accounting file. No errors are reported in the ACMS audit log or in the software event log (SWL) file.

- The account specified for an application does not exist and the application will not start.

Messages reported in prior versions:

```
$ ACMS/START APPL FOO
%ACMSOPR-E-STRTAPLERR, Error while attempting to START APPLICATION FOO
-ACMSACC-E-APPLTERM, EXC process for application FOO terminated
-LOGIN-F-NOSUCHUSER, no such user
%ACMSOPR-E-ERROR, Some operations may not have been performed
```

Messages reported in OpenVMS Version 7.3-2:

```
$ ACMS/START APPL FOO
%ACMSOPR-E-STRTAPLERR, Error while attempting to START APPLICATION FOO
-ACMSACC-E-APPLTERM, EXC process for application FOO terminated
%ACMSOPR-E-ERROR, Some operations may not have been performed
```

The exact error can be seen in the accounting file. The SWL file also has a log of the failure.

2.3 BASIC: V1.5A Required to Create STARLET Library

V7.3-2

Because of a change in OpenVMS Version 7.3-2, BASIC versions prior to V1.5A cannot create the BASIC\$STARLET library file during installation.

Earlier versions of BASIC can install on OpenVMS Version 7.3-2 provided you do not request the option to build the STARLET library file. Also, previously installed BASIC compilers and previously created STARLET library files will continue to function after upgrading an older OpenVMS system to Version 7.3-2.

It is only the BASIC\$STARLET library file creation that does not work on OpenVMS Version 7.3-2. The BASIC V1.5A kit contains an enhanced installation procedure that correctly builds the STARLET library file on OpenVMS Version 7.3-2.

BASIC V1.5A is available on the latest consolidated layered product CD-ROM.

2.4 COBOL: Changes in I/O Run-Time Diagnostics and RMS Special Registers

V7.3

OpenVMS Associated Products Release Notes

2.4 COBOL: Changes in I/O Run-Time Diagnostics and RMS Special Registers

Because of the addition of Extended File Support in OpenVMS Alpha Version 7.2, you may notice changes in the handling of I/O run-time diagnostics and RMS special registers on OpenVMS Alpha Version 7.2 and higher. In particular, a long file name that produced RMS\$_FNM under versions of OpenVMS Alpha prior to Version 7.2 now produces RMS\$_CRE on OpenVMS Alpha Version 7.2 and higher. You do not need to use the new ODS-5 support to see the RMS differences.

2.5 COM for OpenVMS

V7.3-2

For the latest information about COM for OpenVMS, please refer to the following website:

<http://h71000.www7.hp.com/openvms/PRODUCTS/DCOM/INDEX.HTML>

2.5.1 NTA\$LOGON /TYPE Qualifier Behavior Changes

V7.3-2

Because of changes to the \$ACM system service (see Section 5.20), the NTA\$LOGON utility now behaves differently in two ways:

- The NTA\$LOGON /TYPE qualifier no longer requires IMPERSONATE privilege.
- BATCH is no longer a valid value for the NTA\$LOGON /TYPE qualifier. NTA\$LOGON now returns an ACME-E-INVREQUEST error if you explicitly specify /TYPE=BATCH, or if you invoke the NTA\$LOGON utility in batch mode without explicitly specifying the /TYPE qualifier with a value other than BATCH.

To avoid this error, specify the /TYPE=NETWORK qualifier in batch mode. For example:

```
$ NTLOG*ON      := "$NTA$LOGON.EXE"  
$ NTLOGON /TYPE=NETWORK NTUSER1 "password"
```

2.5.2 Registry Access Error with Heavy Load of Applications

V7.3-2

You might get an "Error accessing registry database, contact system manager (0x000025fc)" message if you run a heavy load of COM for OpenVMS applications with the CTLPAGES value set to 256 or less. Set the CTLPAGES value to 512 to avoid this problem.

2.6 DECdfs Version 2.3-3 Required for OpenVMS Version 7.3-2

V7.3-2

To run DECdfs for OpenVMS on OpenVMS Alpha Version 7.3-2, you must install DECdfs Version 2.3-3, which ships with OpenVMS Version 7.3-2. Earlier versions of DECdfs on OpenVMS Alpha Version 7.3-2 will fail to start, resulting in a system bugcheck.

If you have an old version of DECdfs on your system, comment out the @SYS\$STARTUP:DFS\$STARTUP command from your system startup procedure (SYS\$MANAGER:SYSTARTUP_VMS.COM) before upgrading to OpenVMS Version 7.3-2. After you upgrade to OpenVMS Version 7.3-2 and install DECdfs for OpenVMS Version 2.3-3, you can reinstate (uncomment) the DECdfs startup command.

OpenVMS Associated Products Release Notes

2.7 DECforms Web Connector Version 3.0

2.7 DECforms Web Connector Version 3.0

V7.3-1

If you already have DECforms installed, perform the following tasks to enable DECforms Web Connector V3.0 to run on OpenVMS Version 7.3-1 and higher:

1. Remove or comment out the following line:

```
$ @SYS$COMMON: [JAVA$122.COM] JAVA$122_SETUP.COM
```

from these command procedures in the FORMS\$INSTALL_AREA directory:

- FORMS_SMGR_STARTUP.COM
 - FORMS_WEB\$STARTUP.COM
 - FORMS_WEB_CONFIG.COM
2. Ensure that the Java™ environment is set up systemwide for *all* processes. HP recommends adding the Java environment setup to the system's SYLOGIN.COM file.
 3. Ensure that the browser clients use the Sun Java Plugin Version 1.2.2, as stated in the SPD and the Administrative guide.

2.8 DEC PL/I: RTL Support for OpenVMS

V7.3

There is a known incompatibility between the PL/I RTL distributed with the OpenVMS operating system and the more recent PL/I RTL owned and distributed by Kednos Corporation. The older version shipped with the OpenVMS operating system may overwrite a newer version. The image in question is SYS\$LIBRARY:DPLI\$RTL\$SHR.EXE.

OpenVMS distributes the following version of the file, which can be identified by using the DCL command ANALYZE/IMAGE:

Image Identification Information

```
image name: "DPLI$RTL$SHR"  
image file identification: "V4.0-6"
```

If you execute an ANALYZE/IMAGE command before upgrading to OpenVMS Version 7.3 or higher and find a newer version of DPLI\$RTL\$SHR.EXE, you can either copy it and restore it after the upgrade or reinstall the PL/I kit afterward.

Any questions about DEC PL/I and VAX PL/I should be directed to Kednos Corporation:

Phone: (831) 373-7003
Email: tom@kednos.com

2.9 FMS ECO Kit

V7.3-2

To install FMS on OpenVMS Alpha, use the following FMS ECO kit (or a later one):

DECFMSECO5024

2.10 Graphical Configuration Manager (GCM)

The Graphical Configuration Manager (GCM) is included on the Layered Products CD-ROM that ships with the operating system. However, GCM is frequently updated. Check regularly for new versions of the software on the following web page:

<http://h71000.www7.hp.com/openvms/products/gcm/index.html>

2.11 HP DECram

This section contains release notes pertaining to DECram.

2.11.1 DECram Version 3.2 Required for OpenVMS Version 7.3-2

V7.3-2

DECram Version 3.2 is required for OpenVMS Version 7.3-2. This version of DECram ships on the Layered Products CD-ROM and on the Software Products Library CD-ROM.

2.11.2 DECram Version 2.5

V7.3-2

OpenVMS Alpha Version 7.3-2 is the last version to support DECram Version 2.5. In future releases of OpenVMS Alpha, DECram will be included as a System Integrated Product (SIP), and will continue to require a unique license. DECram Version 2.5 will continue to be supported on VAX systems only.

2.11.3 Maximum Disk Size for DECram Version 3.*

V7.3-2

Contrary to what the DECram documentation cites, theoretically the largest value that you can specify for /SIZE and /CAPACITY for DECram Version 3.* is 4,294,967,295 (or %FFFFFFFF), depending on your available free memory.

However, the largest value that OpenVMS supports is 2,147,483,647 for ODS-2 and ODS-5 volumes.

2.11.4 Earlier Versions of DECram

V7.3-2

DECram Version 3.2 runs on OpenVMS Alpha systems, including Galaxy configurations, running OpenVMS Alpha Version 7.2-2 or higher. DECram Version 3.2 and supporting documentation are included on the Software Products Library and the Online Documentation Library CD-ROMs.

DECram Version 2.5 runs on all versions of OpenVMS for both VAX and Alpha systems. (However, OpenVMS Version 7.3-2 is the last release to support DECram Version 2.5. See Section 2.11.2.) Version 2.5, like Version 3.2, uses IOPOST_LOCAL postprocessing and allocates memory across resource affinity domains (RADs).

The following table summarizes OpenVMS support for versions of DECram:

OpenVMS Associated Products Release Notes

2.11 HP DECram

DECram Version	Supported on OpenVMS Alpha Versions	Supported on OpenVMS VAX Versions
V3.2-3HP	V7.2-2 and higher	No
V2.5	All; see Section 2.11.2	All

2.11.5 DECram Commands and Errors

V7.3-1

It is important to check for disk errors after issuing any DECram command, because not all errors are returned to the user interface. Errors specific to a device are sent to the system error log. Type SHOW DEVICE MD at the DCL prompt to see if any device errors were generated as a result of a DECram command. You will need to use an error log analyzer tool to recover the error. Errors are logged in ASCII file format; you can search for errors with an MD-E-FAILURE prefix in the SYS\$SYSROOT:[SYSERR]ERRLOG.SYS file.

2.11.6 DECram and Volume Shadowing

V7.3-1

Using Volume Shadowing for OpenVMS, DECram Version 3.1 and higher can shadow a DECram disk to a physical disk. However, be aware that in the current implementation of Volume Shadowing for OpenVMS, if the physical disk goes away, you will be writing to a volatile disk. A mechanism that will “freeze writes” if the physical disk goes away is planned for a future release of Volume Shadowing for OpenVMS.

2.12 HP DECwindows Motif for OpenVMS

This section contains release notes pertaining to the HP DECwindows Motif for OpenVMS product.

2.12.1 Supported Versions

HP DECwindows Motif for OpenVMS Alpha Version 1.3-1 is the only version of DECwindows Motif supported in the OpenVMS Alpha Version 7.3-2 environment. When you install or upgrade to OpenVMS Alpha Version 7.3-2, the DECwindows Motif Version 1.3-1 software is installed automatically if you answer Yes to the question during installation.

2.12.2 Support for LAT Transport Interface Restored

V7.3-2

With OpenVMS Alpha Version 7.3-2, support for the DECwindows Motif interface to the LAT transport, which was withdrawn with HP DECwindows Motif for OpenVMS Alpha Version 1.3, has been restored. This support enables users of X terminal systems, such as the VXT2000, to start LAT X sessions to communicate with systems running DECwindows Motif Version 1.3-1 or higher. It also allows single- and multithreaded client applications running on these DECwindows Motif systems to use the LAT transport to connect to X terminal systems.

Note that the restored LAT interface included with the OpenVMS Alpha Version 7.3-2 operating system can be used as a valid network transport for communication with the DECwindows Motif Version 1.3 and OpenVMS Alpha Version 7.3-2 display servers. However, use with any other communication

OpenVMS Associated Products Release Notes

2.12 HP DECwindows Motif for OpenVMS

protocols in the X11R6.6 environment is not supported. This includes communication by or with the following:

- Inter-Client Exchange (ICE) and Session Manager protocols
- Low-Bandwidth X (LBX) proxy servers
- Proxy manager applications
- Font servers

Additionally, HP does not support the use of token-based authentication protocols (such as MIT-MAGIC-COOKIE-1 or MIT-KERBEROS-5) with the restored LAT transport interface.

2.12.3 User-Written Transports No Longer Supported

V7.3-2

In DECwindows Motif Version 1.3 for OpenVMS Alpha, significant changes were made to the DECwindows Motif transport library to accommodate multithreading and the communication needs of the Inter-Client Exchange (ICE) protocol, Low-Bandwidth X (LBX) proxy server, and Input Method servers. As a result, HP has discontinued support for user-written network transports on systems running DECwindows Motif Version 1.3 or higher.

All existing transports (DECNET, TCPIP, and LOCAL) remain available and function as expected. However, HP no longer provides support for designing and implementing user-written transports based on the updated transport interface. The *VMS DECwindows Transport Manual* has been archived, and the new libraries are not publicly available.

If you have implemented a custom transport and want to migrate that transport to the DECwindows Motif Version 1.3-1 or greater environment, contact your HP support representative to develop a migration strategy.

2.12.4 Available Language Variants

V7.3-2

The only language variants offered with DECwindows Motif for OpenVMS Alpha Version 1.3-1 are Hebrew and Japanese.

If you require another language variant for DECwindows Motif, contact your HP support representative either to identify a prior version that offers the variant or to discuss options regarding software translation.

2.12.5 System Parameter Values Required for Installation

V7.3-1

The installation procedure for DECwindows Motif for OpenVMS Version 1.3 or higher can fail if the values for the GBLPAGES, FREE_GBLPAGES, and CLISYMTBL system parameters are set too low.

The installation fails with the following error:

```
%SYSTEM-W-NOSUCHFILE, no such file
\sys$library:decw$xlshr.exe\
```

OpenVMS Associated Products Release Notes

2.12 HP DECwindows Motif for OpenVMS

If the installation fails, set these parameters to the minimum values shown in the following table, then reinstall the product.

	GBLPAGES	FREE_GBLPAGES	CLISYMTBL
Alpha	150000	92000	512
VAX	62000	47000	265

2.13 MultiNet Version 4.3

V7.3

Users of Process Software MultiNet Version 4.3 (or earlier) who are also using DCE should install the UCXDRIVER-041_a043 (or higher) ECO from Process Software. MultiNet Version 4.4 includes this fix.

Refer to the following web site to obtain the most recent patch and any dependencies for your version and/or contact Process Software Technical Support:

<http://www.multinet.process.com/eco.html>

2.14 Pascal

The following release notes pertain to HP Pascal.

2.14.1 V5.8A Required to Create STARLET Library

V7.3-2

Because of a change in OpenVMS Version 7.3-2, Pascal versions prior to V5.8A cannot create the STARLET library files during installation.

Earlier versions of Pascal can install on OpenVMS Version 7.3-2 if you answer "NO" to the option to create and install the STARLET library files. Also, previously installed Pascal compilers and previously created STARLET library files will continue to function after upgrading an older OpenVMS system to Version 7.3-2.

It is only the STARLET library creation portion of the Pascal installation that does not work on OpenVMS Version 7.3-2. The Pascal V5.8A kit contains an enhanced installation procedure to correctly build the STARLET library files on OpenVMS Version 7.3-2.

Pascal V5.8A is available on the latest consolidated layered product CD-ROM.

2.14.2 Installing HP Pascal After an Upgrade

V7.3

This note applies to any version of HP Pascal and any version of the OpenVMS operating system.

After upgrading OpenVMS, you should reinstall HP Pascal to produce new versions of STARLET.PAS and other definition files to match the upgraded system.

If you do not reinstall HP Pascal after upgrading OpenVMS, the compiler on your system will still work correctly. However, STARLET.PAS and the other definition files will not contain any new or corrected definitions supplied by the OpenVMS upgrade.

General User Release Notes

This chapter provides information for all users of the OpenVMS operating system. It includes information about commonly used commands and utilities.

For information about new features included in this version of the software, refer to the *HP OpenVMS Alpha Version 7.3-2 New Features and Documentation Overview*.

3.1 OpenVMS Freeware CD-ROMs

V7.3-2

Included in the OpenVMS Version 7.3-2 CD-ROM kit are the OpenVMS Freeware Version 6.0 CD-ROMs. The Freeware CD-ROMs contain free software tools and utilities for creating applications and managing OpenVMS systems.

To mount the Freeware CD-ROMs, insert a CD-ROM into the CD-ROM drive and enter the following commands appropriate to the freeware volume being mounted. For additional information about the freeware, refer to the `FREWARE_README.TXT` files.

Freeware Volume 1:

```
$ MOUNT ddcu: FREEWARE60_1
$ TYPE DISK$FREEWARE60_1:[FREEWARE]FREEWARE_README.TXT
```

Freeware Volume 2:

```
$ MOUNT ddcu: FREEWARE60_2
$ TYPE DISK$FREEWARE60_2:[FREEWARE]FREEWARE_README.TXT
```

In these `MOUNT` commands, the `ddcu:` specification represents the device name of the CD-ROM device on your OpenVMS system.

If you do not know the name of the CD-ROM device on your system, use the following command to find out:

```
$ PIPE SHOW DEV DK/FULL | SEARCH SYS$INPUT RRD
```

Once the appropriate CD-ROM disk is mounted, you can access the kit directories directly using standard DCL commands such as `DIRECTORY`, or you can use the following command to display the Freeware menu for the volume.

```
$ @ddcu:[FREEWARE]FREEWARE_MENU
```

General User Release Notes

3.2 Online Help Topic Name Changes

3.2 Online Help Topic Name Changes

V7.3-2

In OpenVMS Version 7.3, four traditional online help topic names were changed to distinguish them from new, similar topic names. Since then, the OpenVMS Librarian has been modified so that online help can distinguish an independent topic name that matches the beginning of another topic name. Because of this change, these four topic names have been changed back to their original, shorter names. The following table shows the name changes and lists the names of the help topics for related routines.

Topic Name in V7.3—V7.3-1	Topic Name in V7.3-2 and Prior to V7.3	Related Routine Topic Name
BACKUP_Command	BACKUP	BACKUP_API
FDL_Files	FDL	FDL_Routines
MAIL_Command	MAIL	MAIL_Routines
NCS_Command	NCS	NCS_Routines

Note that you must enter the entire name of these shortened help topics to distinguish them from the longer topic names for the related routines. If you do not enter the full topic name, the help Librarian cannot distinguish the shortcut for the command from the related routine. To specify help for the routine, you must enter the beginning of the topic name up through (and including) the underscore character.

3.3 COPY Command Performance

V7.3-2

To improve the performance of copy operations, the default I/O size used by the COPY command has been increased from 64 blocks to 127 blocks. This change results in a significant reduction of both I/O and CPU time required for copy operations to complete.

Users with very low PGFLQUOTA might experience RMS\$_DME errors as RMS allocates its buffers. To overcome this problem, try increasing the size of PGFLQUOTA.

3.4 ATI RADEON 7500 Graphics

Support for ATI RADEON 7500 Graphics is included in this release. Refer to the *HP OpenVMS Alpha Version 7.3-2 New Features and Documentation Overview* manual for a description of the features supported in OpenVMS Version 7.3-2 and read the release notes in Section 6.9.

3.5 HP Secure Web Browser: Increased Memory Required

V7.3-1

If you have an OpenVMS workstation and you are using the HP Secure Web Browser (SWB), based on Mozilla, the minimum memory requirement is 128 MB; however, 256 MB is highly recommended for more robust performance.

System Management Release Notes

This chapter contains information that applies to system maintenance and management, performance management, and networking.

For information about new features included in this version of the software, refer to the *HP OpenVMS Alpha Version 7.3-2 New Features and Documentation Overview*.

4.1 ACME_SERVER has Multiple Channels to RIGHTSLIST.DAT

V7.3-2

The ACME_SERVER process opens a new channel to the RIGHTSLIST.DAT file each time the server is disabled or reenabled and the server processes at least one of the following requests:

```
$ SET SERVER ACME/DISABLE
$ SET SERVER ACME/ENABLE
```

If this operation is performed a certain number of times, the server may reach its FILLM process quota and become unresponsive.

To resolve the problem, stop the server using one of the following commands:

```
$ STOP/ID
$ SET SERVER ACME/ABORT
```

Then restart the server using one of the following commands:

```
$ SYS$STARTUP:ACME$STARTUP.COM      ! for OpenVMS authentication
$ SYS$STARTUP:NTA$STARTUP_NT_ACME.COM ! for NT authentication with
                                     ! COM for OpenVMS applications
```

4.2 AUTOGEN: New Behavior with NEWPARAMS.DAT Files

V7.3-2

AUTOGEN no longer allows layered product kits to provide NEWPARAMS.DAT records that do not include a product name. The most commonly used products that previously did not adhere to this rule are DECwindows and DECnet-Plus. You must install new versions of both products when you install OpenVMS Alpha Version 7.3-2. (See Section 1.9.5.)

AUTOGEN looks for files named SYS\$SYSTEM:NEWPARAMS.DAT;*, which contain SYSGEN parameter modifications specifying the consumption of system resources by layered products. A software installation kit may provide a NEWPARAMS.DAT file instead of telling the system manager to modify MODPARAMS.DAT to accommodate the requirements of the software being installed. For more information, refer to the AUTOGEN chapter in the *HP OpenVMS System Management Utilities Reference Manual*.

4.3 DECdtm Services

This section describes known problems and restrictions associated with using DECdtm services.

4.3.1 DECdtm/XA with Oracle® 8i and 9i (Alpha Only)

V7.3-2

When you are using DECdtm/XA to coordinate transactions with the Oracle 8i/9i XA Compliant Resource Manager (RM), do not use the dynamic registration XA switch (xaoswd). Version 9.0.1.0.0 of the Oracle shareable library that supports dynamic registration does not work. Always use the static registration XA switch (xaosw) to bind the Oracle RM to the DECdtm/XA Veneer.

The DECdtm/XA V2.1 Gateway now has clusterwide transaction recovery support. Transactions from applications that use a clusterwide DECdtm Gateway Domain Log can now be recovered from any single-node failure. Gateway servers running on the remaining cluster nodes can initiate the transaction recovery process on behalf of the failed node.

4.3.2 IPC-E-BCKTRNSFAIL Error Message

V7.3

This note pertains to ACMS users, possibly Rdb users, and anyone else running a user-written application that calls DECdtm to participate in a distributed transaction with a remote system having these characteristics:

- The network connection is HP DECnet-Plus for OpenVMS.
- The nodes are connected using only an IP router.

Users may see the following error returned by DECnet:

IPC-E-BCKTRNSFAIL, failure on the back translate address request

This error is displayed upon a logical connection failure when the remote node name cannot be translated by DECnet-Plus. The error can be triggered when the DECnet-Plus node name for the remote system is not defined in the local DECnet-Plus database and is defined only as ALIAS in the TCP/IP name server for the remote node. For example, node XXYZZY may be defined as follows:

```
20.43.136.54      XXYZZY.ABC.DEF.COM, XXYZZY
```

To avoid this situation, either define the node name in the local DECnet-Plus database or define the logical SYS\$DECDTM_NODE_NAME to be equivalent to one of the following:

- The value of the system parameter SCSNODE
- The DECnet-Plus simple name
- The TCP/IP alias (a six-character node name string such as XXYZZY, as shown in the preceding example)

For other requirements and restrictions, refer to the section about managing DECdtm Services in the *HP OpenVMS System Manager's Manual*.

4.4 ECP Data Collector and Performance Analyzer V5.4D

V7.3-2

Version 5.4D is the recommended version of Enterprise Capacity and Performance (ECP) Data Collector and the Enterprise Capacity and Performance (ECP) Analyzer for OpenVMS Alpha Version 7.3-2. Both the ECP Data Collector and the ECP Performance Analyzer are backward compatible with OpenVMS Version 6.2 and higher.

4.5 EDIT/FDL: Fixing Recommended Bucket Size

V7.3

Prior to OpenVMS Version 7.3, when running EDIT/FDL, the calculated bucket sizes were always rounded up to the closest disk-cluster boundary, with a maximum bucket size of 63. This could cause problems when the disk-cluster size was large, but the “natural” bucket size for the file was small, because the bucket size was rounded up to a much larger value than required. Larger bucket sizes increase record and bucket lock contention, and can seriously impact performance.

OpenVMS Version 7.3 or higher modifies the algorithms for calculating the recommended bucket size to suggest a more reasonable size when the disk cluster is large.

4.6 Error Log Viewer (ELV) Utility

The following release notes pertain to the Error Log Viewer (ELV) utility for OpenVMS.

4.6.1 Translating Error Log Files from OpenVMS Versions 7.2 through 7.3-1

V7.3-2

When you use the ELV TRANSLATE command to examine error log files created on systems running OpenVMS Version 7.2 through 7.3-1, you might see the following set of messages displayed after the translation of certain types of events:

```
%ELV-E-B2TNOTFND, valid bit-to-text translation data not found  
-ELV-W-NODNOTFND, bit-to-text node not found
```

These messages are the result of a minor change in the error log file format between OpenVMS Versions 7.3-1 and 7.3-2, and can be disregarded. The affected events should otherwise be translated correctly.

4.6.2 Using the /PAGE Qualifier with the TRANSLATE Command

V7.3-2

If a message is signaled while you are viewing a report using the /PAGE qualifier with the TRANSLATE command, the display might become corrupted. The workaround for this problem is to refresh the display using Ctrl/W.

If you press Ctrl/Z immediately after a message is signaled, the program abruptly terminates. The workaround for this problem is to scroll past the signaled message before pressing Ctrl/Z.

4.7 Extended File Cache (XFC) Remedial Kits

V7.3-2

The problems that led to the requirement to disable XFC in Version 7.3 were addressed for Version 7.3-1. Issues were corrected that sometimes led to data corruption and system hangs, and XFC performance was also improved. HP recommends that you reenable XFC.

The minimum amount of memory that XFC allocates was increased from about .25 MB to about 3.2 MB; the latter amount is the same as the default for VIOC.

An XFC kit for OpenVMS Version 7.3 is available that contains all the XFC corrections that were included in OpenVMS Version 7.3-1.

If you have an OpenVMS Cluster system that contains earlier versions of OpenVMS Alpha or OpenVMS VAX and you want to use XFC with OpenVMS Version 7.3 or higher, you must install remedial kits on the systems that are running the earlier versions of OpenVMS. For information about the required kits, see Section 4.13.1.

Caution

The remedial kits correct errors in the cache-locking protocol of VIOC, the predecessor to XFC, and allow older versions of the caches to operate safely with the new XFC. Without the functionality in the remedial kits, however, the system or processes might hang.

4.8 External Authentication

This section contains release notes pertaining to external authentication. External authentication is an optional feature introduced in OpenVMS Version 7.1 that enables OpenVMS systems to authenticate designated users with their external user IDs and passwords. For detailed information about using external authentication, refer to the *HP OpenVMS Guide to System Security*.

4.8.1 SET PASSWORD Behavior Within a DECterm Terminal Session

V7.2

A DECterm terminal session does not have access to the external user name used for login and must prompt for one during SET PASSWORD operations. The external user name defaults to the process's OpenVMS user name. If the default is not appropriate (that is, if the external user name and mapped OpenVMS user name are different), you must enter the correct external user name.

The following example shows a SET PASSWORD operation initiated by a user with the external user name JOHN_DOE. The mapped OpenVMS user name is JOHNDOE and is the default used by the SET PASSWORD operation. In this case, the default is incorrect and the actual external user name was specified by the user.

```
$ set password
External user name not known; Specify one (Y/N) [Y]? Y
External user name [JOHNDOE]: JOHN_DOE
Old password:
New password:
Verification:
%SET-I-SNDEXTAUTH, Sending password request to external authenticator
%SET-I-TRYPWDSYNCH, Attempting password synchronization
$
```

4.8.2 No Password Expiration Notification on Workstations

V7.1

In the LAN Manager domain, a user cannot log in once a password expires.

PC users receive notification of impending external user password expiration and can change passwords before they expire. However, when a user logs in from an OpenVMS workstation using external authentication, the login process cannot determine whether the external password is about to expire. Therefore, sites that enforce password expiration and whose users do not primarily use PCs can choose not to use external authentication for workstation users.

4.9 INITIALIZE Command: Incorrect Messages Are Output

V7.3-2

An error causes incorrect messages to be output with the %INIT-I-LIMITCHANGED message, which informs the user that the value specified by INITIALIZE/LIMIT is being overridden by INITIALIZE.

If you get an %INIT-I-LIMITCHANGED message, check the value specified for the /LIMIT qualifier. It should not be less than the physical size of the disk. Because INITIALIZE detected that the value for /LIMIT is too small, INITIALIZE will use a larger value. INITIALIZE did complete successfully and the disk should mount properly. You can ignore the message or you can reenter the command and either specify a larger value for /LIMIT or specify /LIMIT with no value, in which case the maximum expansion size for the disk will be established.

For more information about /LIMIT, see online help or the description of the INITIALIZE/LIMIT command in the *HP OpenVMS DCL Dictionary*.

4.10 Lock Manager: Fast Lock Remastering and PE1

V7.3

The OpenVMS Distributed Lock Manager has a feature called lock remastering. A lock remaster is the process of moving the lock mastership of a resource tree to another node in the cluster. The node that masters a lock tree can process local locking requests much faster because communication is not required with another node in the cluster. Having a lock tree reside on the node doing the most locking operations can improve overall system performance.

Prior to OpenVMS Version 7.3, lock remastering resulted in all nodes sending one message per local lock to the new master. For a very large lock tree, it could require a substantial amount of time to perform the lock remastering operation. During the operation, all application locking to the lock tree is stalled.

Starting with OpenVMS Version 7.3, sending lock data to the new master is done with very large transfers. This is a much more efficient process and results in moving a lock tree from 3 to 20 times faster.

System Management Release Notes

4.10 Lock Manager: Fast Lock Remastering and PE1

Only nodes running Version 7.3 or higher can use large transfers for lock remastering. Remastering between OpenVMS Version 7.3 or higher nodes and prior version nodes still requires sending a single message per lock.

If you currently use the PE1 system parameter to limit the size of lock trees that can be remastered, HP recommends that you either try increasing the value to allow large lock trees to move or try setting the value to zero (0) to allow any size lock tree to move.

4.11 Logical Disk (LD) Utility: Error when Using RMS

V7.3-2

A feature of the Logical Disk (LD) utility can cause problems for end users who are using logical disks. This problem can occur on any version of OpenVMS that uses the LD utility.

The LD utility bypasses the cache for the container file when it is operating on the disk. If RMS is used to read or write to the container file, RMS will have stale data if the LD utility is used to connect to the file and subsequently to the logical disk that is being written.

This problem occurs mainly when the LD utility is used to create disk images that are to be burned onto a CD-ROM.

To work around the problem, execute the following DCL command to turn caching off for any file that will be used as a container file for the LD utility:

```
$ SET FILE/CACHING_ATTRIBUTE=NO_CACHING CONTAINER_FILE.DSK.
```

This DCL command is not executed as part of the CDRECORD.COM command procedure. Therefore, if you reuse a container file for a logical disk created by CDRECORD.COM, be sure to turn off the caching using this command.

4.12 MAIL Utility: Documentation Correction

V7.3-2

In the “Customizing Mail” section of the “Customizing the Operating System” chapter in the *HP OpenVMS System Manager's Manual*, two values for the MAIL\$SYSTEM_FLAGS logical, 8 and 16, are incorrectly marked as VAX only. These values also apply to Alpha systems.

4.13 OpenVMS Cluster Systems

The release notes in this section pertain to OpenVMS Cluster systems.

4.13.1 Patch Kits Needed for Cluster Compatibility

V7.3-2

Before introducing an OpenVMS Version 7.3-2 system into an existing OpenVMS Cluster system, you must apply certain patch kits (also known as remedial kits) to your systems running earlier versions of OpenVMS. If you are using Fibre Channel, XFC, or Volume Shadowing, additional patch kits are required. Note that these kits are version specific.

Table 4–1 lists the facilities that require patch kits and the patch kit names. Each patch kit has a corresponding readme file with the same name (file extension is .README).

System Management Release Notes

4.13 OpenVMS Cluster Systems

You can either download the patch kits from the following web site (select the OpenVMS Patch Kits option), or contact your HP support representative to receive the patch kits on media appropriate for your system:

<http://h18007.www1.hp.com/support/files/index.html>

Note

Patch kits are periodically updated on an as-needed basis. Always use the most recent patch kit for the facility, as indicated by the version number in the kit's readme file. The most recent version of each kit is the version posted to the web site.

Table 4–1 Patch Kits Required for Cluster Compatibility

Facility	File Name
OpenVMS Alpha Version 7.3-1	
Update kit with all patch kits except DECnet-Plus	DEC-AXPVMS-VMS731_UPDATE-V0100
DECnet-Plus	DEC-AXPVMS-DNVOSIECO02-V0703-1-4.PCSI
OpenVMS Alpha Version 7.3	
Update kit with all patch kits except DECnet-Plus	DEC-AXPVMS-VMS73_UPDATE-V0200–4.PCSI
DECnet-Plus	DEC-AXPVMS-DNVOSIECO03-V0703–4.PCSI
OpenVMS VAX Version 7.3	
Audit Server	VAXAUDS01_073
Cluster	VAXSYSL01_073
DECnet-Plus	VAXVMS-DNVOSIECO02-V0703–4.PCSI
DECwindows Motif	VAXDWMOTMUP01_073
Files 11	VAXF11X02_073
MAIL	VAXMAIL01_073
MOUNT	VAXMOUN01_073
RMS	VAXRMS01_073
Shadowing	VAXSHAD01_073
System	VAXSYS01_073
OpenVMS Alpha Version 7.2-2	
Update kit with all patch kits except those listed below	DEC-AXPVMS-VMS722_UPDATE-V0100-4.PCSI
Cluster	DEC-AXPVMS-VMS722_DRIVER-V0300–4.PCSI
DECnet-Plus	DEC-AXPVMS-DNVOSIECO06-V0702–4.PCSI

(continued on next page)

System Management Release Notes

4.13 OpenVMS Cluster Systems

Table 4–1 (Cont.) Patch Kits Required for Cluster Compatibility

Facility	File Name
OpenVMS VAX Version 7.2	
Update kit with all patch kits except those listed below	VAXUPDATE01_072
Audit Server	VAXAUDS01_072
Backup utility	VAXBACK02_072
CLI Utility	VAXCLIU03_072
C RTL	VAXACRT02_072
DCE	DEC-VAXVMS-VAX_DCEECO_015_1-V0100–4.PCSI
DECnet-Plus	DEC-VAXVMS-DNVOSIECO06-V0702–4.PCSI
DECwindows Motif	VAXDWMOTMUP01_072
Fibre Channel	VAXDRIV02_072
Files 11	VAXF11X04_072
LAT	VAXLAT01_072
LIBRTL	VAXLIBR01_072
MANAGE	VAXMANA01_072
MIME	VAXMIME02_072
Mutex Release Error	VAXDUP01_072
ODS1	VAXODS1_01_072
PCSI	DEC-VAXVMS-VMS72_PCSI-V0101–4.PCSI
PThreads	VAXPTHR01_072
RMS	VAXRMS02_072
Volume Shadowing	VAXSHAD03_072
XFC/VCC compatibility support and logical names	VAXSYS03_072

4.13.2 New API Can Correct Incompatibility of Fibre Channel and SCSI Multipath with Some Third-Party Products

V7.3-2

OpenVMS Alpha Version 7.2-1 introduced the multipath feature, which provides support for failover between the multiple paths that can exist between a system and a SCSI or Fibre Channel device. OpenVMS Alpha Version 7.3-1 introduced support for failover between Fibre Channel multipath tape devices.

This multipath feature can be incompatible with some third-party disk-caching, disk-shadowing, or similar products. HP advises that you do not use such software on SCSI or Fibre Channel devices that are configured for multipath failover until this feature is supported by the producer of the software.

Third-party products that rely on altering the Driver Dispatch Table (DDT) of either the OpenVMS Alpha SCSI disk class driver (SYS\$DKDRIVER.EXE), the OpenVMS Alpha SCSI tape class driver (SYS\$MKDRIVER.EXE), or the SCSI generic class driver (SYS\$GKDRIVER) may need to be modified in order to function correctly with the SCSI multipath feature.

Producers of such software can now modify their software using new DDT Intercept Establisher routines introduced in OpenVMS Alpha Version 7.3-2. For more information about these routines, refer to the *HP OpenVMS Alpha Version 7.3-2 New Features and Documentation Overview* manual.

Note

If you are using a third-party disk-caching product or disk shadowing application, refrain from using it in an OpenVMS SCSI or Fibre Channel multipath configuration until you confirm that the application has been revised using these new routines.

For more information about OpenVMS Alpha SCSI and Fibre Channel multipath features, refer to *Guidelines for OpenVMS Cluster Configurations*.

4.13.3 SCSI Tape Drives: MEDOFL Errors After Tape Dismount

V7.3-2

Occasionally, a “%SYSTEM-F-MEDOFL, medium is offline” error can occur on the first command executed on a SCSI tape after the tape has been dismounted using a DISMOUNT/NOUNLOAD command. For example, this can happen when you attempt to initialize or mount the tape immediately after dismounting it. The error is returned because the tape is still rewinding as part of the dismount operation.

If the tape unit is a member of a multipath set, a path switch might occur (instead of a MEDOFL error) as part of multipath recovery. These MEDOFL errors and path switches are more likely to occur on certain models of SCSI tape drives such as the LTO-2 HP Ultrium 460.

If a path switch occurs on the first command after a tape DISMOUNT command, this indicates that the tape has recovered and no user action is required. If a MEDOFL error occurs, retry the failed command after the tape finishes rewinding. The SCSI tape driver will be modified in a future remedial kit to eliminate the need for such manual retries.

4.13.4 CLUSTER_CONFIG.COM and Limits on Root Directory Names

V7.3-2

This note updates Table 8-3 (Data Requested by CLUSTER_CONFIG_LAN.COM and CLUSTER_CONFIG.COM) in the *OpenVMS Cluster Systems* manual.

The documentation specifies a limit on the number of hexadecimal digits you can use for computers with direct access to the system disk. The limit is correct for VAX computers but not for Alpha computers.

The command procedure prompts for the following information:

Computer's root directory name on cluster system disk:

The documentation currently states:

Press Return to accept the procedure-supplied default, or specify a name in the form SYSx:

- For computers with direct access to the system disk, x is a hexadecimal digit in the range of 1 through 9 or A through D (for example, SYS1 OR SYSA)

System Management Release Notes

4.13 OpenVMS Cluster Systems

- For satellites, *x* must be in the range of 10 through FFFF

The limit on the range of hexadecimal values with direct access to the system disk is correct for VAX computers. For Alpha computers with direct access to the system disk, the valid range of hexadecimal values is much larger. It includes both the VAX range of 1 through 9 or A through D, and also the range 10 through FFFF. Note that SYSE and SYSF are reserved for system use.

The *OpenVMS Cluster Systems* manual will include this information in its next revision.

4.13.5 Booting Satellites Over FDDI in a Mixed-Version Cluster

V7.3

Changes to OpenVMS Version 7.3 (or higher) may affect satellite booting over FDDI for satellites running versions of OpenVMS earlier than Version 7.3. The problem can occur when the system parameter NISCS_LAN_OVRHD is set to a value less than 6 (the default is 18), and the system parameter NISCS_MAX_PKTSZ is set for maximum size FDDI packets (4468). NISCS_LAN_OVRHD decreases the maximum packet size used for LAN communications to accommodate devices such as the DESNC (an Ethernet encryption device). For OpenVMS Version 7.3 or higher, NISCS_LAN_OVRHD is not used, so the maximum packet size is not reduced.

The problem is that the buffer size used by the FDDI boot driver is 12 bytes too short. The FDDI boot driver portion of the satellite boot typically causes 12 bytes of incorrect data (often zeros) to be interspersed throughout the images loaded during SYSBOOT. This generally results in an obscure failure or halt very early in the life of the system (measured in seconds).

The solution is to obtain a Boot Driver patch kit that corrects the problem and to install the patch on the satellite system root. Alternatively, on the systems serving the system disk to the satellite, ensure that the value of the system parameter NISCS_MAX_PKTSZ is at least 12 bytes less than the maximum FDDI packet size.

The following systems are affected:

- Alpha satellite using an FDDI adapter that is booting from an OpenVMS Version 7.3 or higher Alpha or VAX system whose NISCS_MAX_PKTSZ value is greater than 4456.
- Alpha satellite using an FDDI adapter that is booting from a pre-OpenVMS Version 7.3 system, which is serving a system disk via FDDI, and the value of NISCS_MAX_PKTSZ minus NISCS_LAN_OVRHD is greater than 4456. The served system disk may be running OpenVMS Version 7.3 or higher, or an earlier version. The problem is more likely to occur if the system disk is Version 7.3 or higher, because NISCS_LAN_OVRHD is most likely set to 18 for prior versions.

4.13.6 PEdriver Error Message Change

V7.3-2

In the final build of OpenVMS Version 7.3-2, it was discovered that a last-minute bug changed the way the error message is displayed when PEdriver is closing a virtual circuit. Prior to Version 7.3-2, the error message displayed the remote node name. For example:

```
%PEA0, Software is Closing Virtual Circuit - REMOTE NODE LARRY
```

The Version 7.3-2 message displays PEdriver's internally assigned number for the remote port instead of the remote node name. For example:

```
%PEA0, Software is Closing Virtual Circuit - REMOTE PORT 219
```

Unfortunately, there is no easy way to determine the mapping between remote port numbers and the name of the node associated with that numeric value.

This problem will be fixed in the next release.

4.13.7 PEdriver Channels with Priority of -128 Not Used

V7.3-2

Starting with OpenVMS Version 7.3-2, a PEdriver channel whose priority is -128 will not be used for cluster communications. Therefore, you can disable cluster communications for a particular channel by using SCACP or the Availability Manager to set the channel's priority to -128.

A channel's priority is the sum of the management priorities assigned to the local LAN device and the channel itself. Therefore, you can assign any combination of channel and LAN device management priority values to achieve a total of -128.

4.13.8 Cluster Performance Reduced with CI-LAN Circuit Switching

V7.3-1

In rare cases, in an OpenVMS Cluster configuration with both CI and multiple FDDI, 100 Mb/s or Gb/s Ethernet-based circuits, you might observe that SCS connections are moving between CI and LAN circuits at intervals of approximately 1 minute. This frequent circuit switching can result in reduced cluster performance and may trigger mount verification of shadow set members.

PEdriver can detect and respond to LAN congestion that persists for a few seconds. When it detects a significant delay increase or packet losses on a LAN path, the PEdriver removes the path from use. When it detects that the path has improved, it begins using it again.

Under marginal conditions, the additional load on a LAN path resulting from its use for cluster traffic may cause its delay or packet losses to increase beyond acceptable limits. When the cluster load is removed, the path might appear to be sufficiently improved so that it will again come into use.

If a marginal LAN path's contribution to the LAN circuit's load class increases the circuit's load class above the CI's load class value of 140 when the marginal path is included (and, conversely, decreases the LAN circuit's load class below 140 when the path is excluded), SCS connections will move between CI and LAN circuits.

You can observe connections moving between LAN and CI circuits by using SHOW CLUSTER with the CONNECTION and CIRCUITS classes added.

Workarounds

If excessively frequent connection moves are observed, you can use one of the following workarounds:

- You can use SCACP or AM to assign a higher priority to the circuit, or the port you wish to be used, thus overriding automatic connection assignment and moving.

System Management Release Notes

4.13 OpenVMS Cluster Systems

Examples of SCACP commands are:

```
$ MC SCACP
SCACP> SET PORT PNA0 /PRIORITY=2      ! This will cause circuits from local
                                       ! CI port PNA0 to be chosen over
                                       ! lower priority circuits.
```

```
SCACP> SET PORT PEA0 /PRIORITY=2      ! This will cause LAN circuits to be
                                       ! chosen over lower priority circuits.
```

- You can use the SCACP SHOW CHANNEL commands to determine which channels are being switched into/out of use. Then you can use SCACP to explicitly exclude a specific channel by assigning it a lower priority value than the desired channels. For example:

```
SCACP> SET CHANNEL LARRY /LOCAL=EWB/REMOTE=EWB /PRIORITY=-2
```

Note that CHANNEL and LAN device priority values in the range of *max*, *max-1* are considered equivalent; that is, they are treated as if they both had the maximum priority value. A difference of 2 or more in priority values is necessary to exclude a channel or LAN device from use.

4.13.9 Gigabit Ethernet Switch Restriction in an OpenVMS Cluster System

Permanent Restriction

Attempts to add a Gigabit Ethernet node to an OpenVMS Cluster system over a Gigabit Ethernet switch will fail if the switch does not support autonegotiation. The DEGPA enables autonegotiation by default, but not all Gigabit Ethernet switches support autonegotiation.

In addition, the messages that are displayed may be misleading. If the node is being added using CLUSTER_CONFIG.COM and the option to install a local page and swap disk is selected, the problem may look like a disk-serving problem. The node running CLUSTER_CONFIG.COM displays the message “waiting for *node-name* to boot,” while the booting node displays “waiting to tune system.” The list of available disks is never displayed because of a missing network path. The network path is missing because of the autonegotiation mismatch between the DEGPA and the switch.

To avoid this problem, disable autonegotiation on the new node's DEGPA, as follows:

- Perform a conversational boot when first booting the node into the cluster.
- Set the new node's system parameter LAN_FLAGS to a value of 32 to disable autonegotiation on the DEGPA.

4.13.10 Multipath Tape Failover Restriction

V7.3-1

While the INITIALIZE command is in progress on a device in a Fibre Channel multipath tape set, multipath failover to another member of the set is not supported. If the current path fails while another multipath tape device is being initialized, retry the INITIALIZE command after the tape device fails over to a functioning path.

This restriction will be removed in a future release.

4.13.11 No Automatic Failover for SCSI Multipath Medium Changers

V7.3-1

Automatic path switching is not implemented in OpenVMS Alpha Version 7.3-1 or higher for SCSI medium changers (tape robots) attached to Fibre Channel using a Fibre-to-SCSI tape bridge. Multiple paths can be configured for such devices, but the only way to switch from one path to another is to use manual path switching with the SET DEVICE/SWITCH command.

This restriction will be removed in a future release.

4.14 OpenVMS Galaxy

OpenVMS provides Galaxy support on AlphaServer ES47, ES80, and GS1280 systems. Galaxy support on these systems requires Version 6.6 firmware and may require additional Version 7.3-2 patch kits. The firmware can be obtained from the following website:

<http://ftp.digital.com/pub/Digital/Alpha/firmware/interim/gs1280/>

Eventually, the Version 6.6 firmware will also be available on CD-ROM.

The following sections contain release notes pertaining to OpenVMS Galaxy systems. Also see related notes in Section 6.5.

4.14.1 OpenVMS Graphical Configuration Manager

The OpenVMS Graphical Configuration Manager (GCM) is not supported for AlphaServer ES47/ES80/GS1280 Galaxy configurations at this time. However, the Graphical Configuration Utility (GCU) is supported. This restriction will be removed in the future.

4.14.2 Smart Array 5300 Restrictions

The Smart Array 5300 (KZPDC) Backplane Raid Controller is currently supported only as a data device in ES47/ES80/GS1280 Galaxy configurations. Boot and crash dump capability are not supported at this time on these controllers. The goal is to provide support with corrected firmware or corrected OpenVMS software.

For information about configuring a Galaxy on an AlphaServer ES47/ES80/GS1280 system, see the *HP OpenVMS Alpha Partitioning and Galaxy Guide*.

4.14.3 Firmware and Patch Kit Requirements

Hard partition support, which requires a firmware update and a patch kit, has been qualified and is now available on the AlphaServer ES47/ES80/GS1280 systems. The *HP OpenVMS Alpha Partitioning and Galaxy Guide* provides more information about the firmware and patch kit requirements and describes how to configure hard partitions on these systems.

Note

The former limitation of hard partitions on system building block boundaries, only, has been removed. Hard partitions on subsystem building block boundaries are now supported, as described in the *HP OpenVMS Alpha Partitioning and Galaxy Guide*. Please note the constraints on hard partitions in subsystem building blocks as described

System Management Release Notes

4.14 OpenVMS Galaxy

in the *HP OpenVMS Alpha Partitioning and Galaxy Guide*. Hard partitions on ES47/ES80/GS1280 systems can support up to 64 processors.

4.14.4 Shared-Memory Global Section Creation Can Return Incorrect Status

V7.3-2

Calls to SYSSCRMPSC_GDZRO_64 with the flag SEC\$M_SHMGS can fail with status SS\$_INFMEM instead of status SS\$_INSF_SHM_REG.

The most likely explanation for this error is that the Galaxy shared-memory code has run out of internal SHM_REG data structures. To correct this condition, increase the value of the SYSGEN parameter GLX_SHM_REG and reboot all Galaxy instances with this larger parameter value.

Note that each SHM_REG data structure consumes only a small amount of memory. Therefore, you can safely increase this parameter to a relatively high number (for example, double the number of expected shared-memory regions) to avoid changing this parameter in small increments and having to reboot the entire Galaxy more than once.

In a mixed-version cluster, driver kits VMS73_DRIVER-V0300 or later and VMS722_DRIVER-V0300 or later should be installed to avoid Galaxy shared-memory interconnect errors.

4.14.5 Galaxy on ES40: Uncompressed Dump Limitation

Permanent Restriction

On AlphaServer ES40 Galaxy systems, you cannot write a raw (uncompressed) dump from instance 1 if instance 1's memory starts at or above 4 GB (physical). Instead, you must write a compressed dump.

4.14.6 Galaxy on ES40: Turning Off Fast Path

V7.3-1

When you implement Galaxy on an AlphaServer ES40 system, you must turn off Fast Path on instance 1. Do this by setting the SYSGEN parameter FAST_PATH to 0 on that instance.

If you do not turn off Fast Path on instance 1, I/O on instance 1 will hang when instance 0 is rebooted. This hang will continue until the PCI bus is reset and instance 1 rebooted. If there is shared SCSI or Fibre Channel, I/O will hang on the sharing nodes and all paths to those devices will be disabled.

4.15 OpenVMS Management Station

V7.3-2

Version 3.2B is the recommended version of OpenVMS Management Station for OpenVMS Alpha Version 7.3-2. However, OpenVMS Management Station is backward compatible with OpenVMS Version 6.2 and higher.

The OpenVMS Alpha Version 7.3-2 installation includes OpenVMS Management Station Version 3.2B, which is also available on the web.

4.16 OpenVMS Registry Can Corrupt Version 2 Format Database

V7.3-2

If you create eight or more volatile subkeys in a key tree and then reboot a standalone system or a cluster, the OpenVMS Registry server can corrupt a Version 2 format Registry database when the server starts up after the reboot.

To avoid this problem, do one of the following:

- Do not use volatile keys.
- Use a Version 1 format database.

Note that Advanced Server for OpenVMS and COM for OpenVMS do not create volatile keys.

4.17 RMS Journaling

The following release notes pertain to RMS Journaling for OpenVMS.

For more information about RMS Journaling, refer to the *RMS Journaling for OpenVMS Manual*. You can access this manual on the OpenVMS Documentation CD-ROM (in the archived manuals directory).

4.17.1 Recovery Unit Journaling Incompatible with Kernel Threads

V7.3

Because DECdtm Services is not supported in a multiple kernel threads environment and RMS recovery unit journaling relies on DECdtm Services, RMS recovery unit journaling is not supported in a process with multiple kernel threads enabled.

4.17.2 Modified Journal File Creation

V7.2

Prior to Version 7.2, recovery unit (RU) journals were created temporarily in the [SYSJNL] directory on the same volume as the file that was being journaled. The file name for the recovery unit journal had the form RMS\$*process_id* (where *process_id* is the hexadecimal representation of the process ID) and a file type of RMS\$JOURNAL.

The following changes have been introduced to RU journal file creation in OpenVMS Version 7.2:

- The files are created in node-specific subdirectories of the [SYSJNL] directory.
- The file name for the recovery unit journal has been shortened to the form: *YYYYYYYY*, where *YYYYYYYY* is the hexadecimal representation of the process ID in reverse order.

These changes reduce the directory overhead associated with journal file creation and deletion.

The following example shows both the previous and current versions of journal file creation:

Previous versions: [SYSJNL]RMS\$214003BC.RMS\$JOURNAL;1
Current version: [SYSJNL.NODE1]CB300412;1

If RMS does not find either the [SYSJNL] directory or the node-specific directory, RMS creates them automatically.

System Management Release Notes

4.17 RMS Journaling

4.17.3 Remote Access of Recovery Unit Journalled Files in an OSI Environment

V6.1

OSI nodes that host recovery unit journalled files that are to be accessed remotely from other nodes in the network must define SYSS\$NODE to be a Phase IV-style node name. The node name specified by SYSS\$NODE must be known to any remote node attempting to access the recovery unit journalled files on the host node. It must also be sufficiently unique for the remote node to use this node name to establish a DECnet connection to the host node. This restriction applies only to recovery unit journalled files accessed across the network in an OSI or mixed OSI and non-OSI environment.

4.17.4 After-Image (AI) Journaling

V6.0

You can use after-image (AI) journaling to recover a data file that becomes unusable or inaccessible. AI recovery uses the AI journal file to roll forward a backup copy of the data file to produce a new copy of the data file at the point of failure.

In the case of either a process deletion or system failure, an update can be written to the AI journal file, but not make it to the data file. If only AI journaling is in use, the data file and journal are not automatically made consistent. If additional updates are made to the data file and are recorded in the AI journal, a subsequent roll forward operation could produce an inconsistent data file.

If you use Recovery Unit (RU) journaling with AI journaling, the automatic transaction recovery restores consistency between the AI journal and the data file.

Under some circumstances, an application that uses only AI journaling can take proactive measures to guard against data inconsistencies after process deletions or system failures. For example, a manual roll forward of AI-journalled files ensures consistency after a system failure involving either an unshared AI application (single accessor) or a shared AI application executing on a standalone system.

However, in a shared AI application, there may be nothing to prevent further operations from being executed against a data file that is out of synchronization with the AI journal file after a process deletion or system failure in a cluster. Under these circumstances, consistency among the data files and the AI journal file can be provided by using a combination of AI and RU journaling.

4.17.5 VFC Format Sequential Files

VAX V5.0

Alpha V1.0

You cannot update variable fixed-length control (VFC) sequential files when using before-image or recovery unit journaling. The VFC sequential file format is indicated by the symbolic value FAB\$_VFC in the FAB\$_RFM field of the FAB.

4.18 Security: Changes to DIRECTORY Command Output

V7.3-2

In OpenVMS Version 7.1 and higher, if you execute the DCL command `DIRECTORY/SECURITY` or `DIRECTORY/FULL` for files that contain Advanced Server (PATHWORKS) access control entries (ACEs), the hexadecimal representation for each Advanced Server ACE is no longer displayed. Instead, the total number of Advanced Server ACEs encountered for each file is summarized in the message, "Suppressed *n* PATHWORKS ACEs."

To display the suppressed ACEs, use the `SHOW SECURITY` command. You must have the `SECURITY` privilege to display these ACEs. Note that, in actuality, the command displays OpenVMS ACEs, including the %x86 ACE that reveals the Windows NT® security descriptor information. The Windows NT security descriptor information pertains to the Advanced Server.

4.19 Server Management Process (SMHANDLER)

V7.3-2

The server management process, SMHANDLER, now starts automatically on Alpha systems that support it. System managers should remove references to the obsolete startup file, `SYS$STARTUP:SYS$SMHANDLER_STARTUP.COM`, from `SYSTARTUP_VMS.COM` or other site-specific startup files. This reference has been removed from `SYSTARTUP_VMS.TEMPLATE`.

Background: What is SMHANDLER?

On certain Alpha systems, the server management process is started to assist the system firmware in reporting and responding to imminent hardware failures. Failure conditions vary but typically include over-temperature conditions, fan failures, or power supply failures. SMHANDLER may report warning conditions to OPCOM, and may initiate a shutdown of OpenVMS if system firmware is about to power off a failing system. In most situations, a controlled shutdown of OpenVMS would be less disruptive than abrupt loss of system power.

To ensure the longest possible up time, system managers can set the `POWEROFF` system parameter to 0. This prevents SMHANDLER from shutting down OpenVMS on a failing system but does not prevent system firmware from powering off the system.

4.20 SYSGEN: Security Auditing Fixed

V7.3-2

Previously, enabling SYSGEN audits, or alarms, did not provide any audits or alarms with information about the parameters being modified. As of OpenVMS Version 7.3-2, this problem is corrected. Audits or alarms now provide a list of the changed parameters along with their old and new values.

4.21 SYSMAN Utility

The following release notes pertain to the SYSMAN utility.

4.21.1 DUMP_PRIORITY LOAD Command

V7.3-2

When you use the SYSMAN command `DUMP_PRIORITY LOAD` to load the contents of the System Dump Priority registry into memory, any entries with a UIC specification of `[group-id,*]` (for example, `[TCPIP$AUX,*]`) are ignored — even though the specification is correct. This will be fixed in a remedial kit for OpenVMS Alpha Version 7.3-2.

Workaround:

You can work around this problem by using the SYSMAN command `DUMP_PRIORITY MODIFY` to remove `,"*` from the System Dump Priority registry entry. For example:

```
SYSMAN> DUMP_PRIORITY MODIFY TCPIP$*/UIC=[TCPIP$AUX,*] /NEWUIC=[TCPIP$AUX]
```

This workaround creates the identical in-memory contents during a `DUMP_PRIORITY LOAD` command. Any subsequent `DUMP_PRIORITY SHOW` command displays it as `[group-id,*]`.

The workaround will continue to work even after the remedial kit is installed.

4.21.2 XP LUNs Can Fail to Autoconfigure

V7.3-2

The SYSMAN IO `AUTOCONFIGURE` command does not configure LUNs on an XP port when a Fibre Channel adapter on a booted host is given access to those LUNs. This problem occurs because of the way the XP array implements selective storage presentation.

You could configure the LUNs by rebooting the host, but any of the following safe, less disruptive workarounds will do the job and can be performed prior to executing SYSMAN IO `AUTOCONFIGURE`:

- Reset the switch port to which the XP port is connected.
Depending on the switch type and the access available to it, this might require executing `portDisable` and `portEnable` commands or by using a web-based GUI to reset the link.
- Remove and replace either end of the Fibre Channel cable between the XP array and the switch port to which it's connected.
- Generate a link event on each host Fibre Channel adapter that has been given access to the LUNs. Do this by running `SY$ETC:FC$CP` for the FGx0 device in the following suggested sequence:
 1. Run `FC$CP` once for the adapter to see what the current "enables" value is set to.
 2. Run `FC$CP` again to set a different value for "enables". (The application will not reset the link unless necessary — and it is not necessary unless the value of "enables" is changed.)
 3. Run `FC$CP` a third time to restore the original value of "enables".

The first two options have the benefit of being seen by all host Fibre Channel adapters on the fabric that have been given access to the LUNs. However, these options also briefly block access to all other LUNs on the XP port and cause mounted volumes to go through Mount Verify.

The last option limits the disruption to those LUNs whose current path is through the affected Fibre Channel adapter — even those that are not on the affected XP port — but the process must be performed for individual adapters.

4.22 System Parameters

The following sections contain notes related to system parameters.

4.22.1 New System Parameters

V7.3-2

The following system parameters are new in OpenVMS Alpha Version 7.3-2:

MVSUPMSG_INTVL
MVSUPMSG_NUM
RMS_CONPOLICY
SHADOW_REC_DLY
SHADOW_SITE_ID

Refer to online help for the definitions of these new parameters.

4.22.2 Modified System Parameters

V7.3-2

Definitions of the following system parameters have been modified in OpenVMS Alpha Version 7.3-2:

CLUSTER_CREDITS
CRD_CONTROL
DUMPSTYLE
FAST_PATH_PORTS
GH_RES_DATA
GH_RSRVPGCNT
LAN_FLAGS
MSCP_CREDITS
MSCP_LOAD
NPAG_BAP_*
RSRVPAGCNT
SCH_CTLFLAGS
SCSNODE
SECURITY_POLICY
SHADOW_MAX_UNIT
TTY_RSPEED
TTY_SPEED

Refer to online help for changes in the definitions of these parameters.

System Management Release Notes

4.22 System Parameters

4.22.3 Obsolete System Parameters

V7.3-2

The following system parameters have been removed in OpenVMS Alpha Version 7.2-3:

SHADOW_REMOVE_1
SHADOW_REMOVE_2

4.22.4 SCSNODE Parameter Size Now Strictly Enforced

V7.3-2

Previously, the documented maximum size of 6 characters for the SYSGEN parameter SCSNODE was not strictly enforced. You could set SCSNODE to a name containing 8 characters, which could lead to device names being too long and producing undesirable effects.

The documented maximum of 6 characters is now strictly enforced. The value of SCSNODE will be truncated by SYSBOOT if the size is set to more than 6 characters in the system parameter file.

4.22.5 AlphaServer GS Series: NPAGERAD System Parameter Default Behavior

V7.3-1

On AlphaServer GS series processors on OpenVMS systems prior to Version 7.3-1, system managers frequently saw pool expansion that increasing NPAGEDYN did not reduce. This problem was caused by leaving NPAGERAD at its default value of 0.

Starting with OpenVMS Version 7.3-1, when NPAGERAD is 0 (the default), the system calculates a value to use for NPAGERAD with the following formula:

$$\text{NPAGEDYN} * \left(1 - \frac{\text{Base RAD memory}}{\text{Total memory}}\right)$$

This calculation gives more pool to the nonbase RADs than before and so reduces the expansion of nonbase RAD pool.

4.23 Terminal Fallback Facility (TFF)

On OpenVMS Alpha systems, the Terminal Fallback Facility (TFF) includes a fallback driver (SYS\$FBDRIVER.EXE), a shareable image (TFFSHR.EXE), a terminal fallback utility (TFU.EXE), and a fallback table library (TFF\$MASTER.DAT).

Note

TFFSHR has been removed from IMAGELIB because it is not a documented, user-callable interface. The image is still available in the SYS\$LIBRARY: directory.

To start TFF, invoke the TFF startup command procedure located in SYS\$MANAGER, as follows:

```
$ @SYS$MANAGER:TFF$SYSTARTUP.COM
```

System Management Release Notes

4.23 Terminal Fallback Facility (TFF)

To enable fallback or to change fallback characteristics, invoke the Terminal Fallback Utility (TFU), as follows:

```
$ RUN SYS$SYSTEM:TFU
TFU>
```

To enable default fallback to the terminal, enter the following DCL command:

```
$ SET TERMINAL/FALLBACK
```

OpenVMS Alpha TFF differs from OpenVMS VAX TFF in the following ways:

- On Alpha systems, the TFF fallback driver is named SYS\$FBDRIVER.EXE. On VAX systems, the TFF fallback driver is named FBDRIVER.EXE.
- On Alpha systems, TFF is capable of handling 16-bit character fallback. The OpenVMS Alpha fallback table library (TFF\$MASTER.DAT) contains four more 16-bit character tables than the VAX library. Table 4–2 describes these additional tables.

Table 4–2 TFF Character Fallback Tables

Table Name	Base	Description
BIG5_HANYU	BIG5	BIG5 for CNS 11643 (SICGCC) terminal/printer
HANYU_BIG5	CNS	CNS 11643 (SICGCC) for BIG5 terminal/printer
HANYU_TELEX	CNS	CNS 11643 for MITAC TELEX-CODE terminal
HANGUL_DS	KS	KS for DOOSAN 200 terminal

These tables are used mainly by the Asian region. Also, the table format was changed due to the support of 16-bit character fallback.

- On Alpha systems, the TFU command SHOW STATISTICS does not display the size of the fallback driver (SYS\$FBDRIVER.EXE).

RT terminals are not supported by TFF.

For more information about the Terminal Fallback Facility, refer to the *OpenVMS Terminal Fallback Utility Manual*. You can access this manual on the OpenVMS Documentation CD-ROM (in the archived manuals directory).

4.24 UETP: Error With DEGPA Device

V7.3-2

The User Environment Test Package (UETP) does not currently support DEGPA network devices. If you get an unsupported device error when running UETP and you have a DEGPA device, edit the UETPSUPDEV.DAT file in the SYSTEST directory, and add the following line just before the last line in the file:

```
20 5D UETUNAS00.EXE ! DEGPA
```

Then rerun UETP. The device should now test properly.

4.25 VMSINSTAL: NOBROADCAST Option

V7.3-2

Previously, if HELPLIB was in use during a VMSINSTAL installation, messages were automatically sent to all users. VMSINSTAL has been modified to check for logical name VMSINSTAL_NOBROADCAST. If the logical name is set, VMSINSTAL will not send a reply to all users when HELPLIB is in use.

4.26 Volume Shadowing for OpenVMS

The following release notes pertain to HP Volume Shadowing for OpenVMS, also known as host-based volume shadowing (HBVS).

4.26.1 Device Name Requirement

V7.3-2

Volume Shadowing for OpenVMS supports device names whose *ddc* portion of the full device name of *\$alloclass\$ddcu*: is three characters.

Prior to this release, it was possible to create device names whose *ddc* portion of the full device name was longer, such as \$1\$DECRAM10:, and these devices mounted successfully. However, mounting such devices as part of a shadow set caused operational problems, such as %MOUNT-F-XSMBRS errors when other disks were added to the shadow set.

Starting with OpenVMS Alpha Version 7.3-2, the Mount utility enforces this three-character requirement for the *ddc* portion of the full device name during the initial attempt to mount the device. If you attempt to mount a device whose name does not conform to this requirement, the following error message is displayed:

```
MOUNT-F-NOTSHDWDEV, not a valid shadow set member
```

4.26.2 Warning About Using SET SHADOW and SHOW SHADOW in DCL Command Procedures

V7.3-2

The new DCL commands SET SHADOW and SHOW SHADOW will continue to evolve. In a future release, the default display and implementation of a SHOW SHADOW/FULL display will change the current formatting. Therefore, HP advises customers not to rely on parsing the current format of output in DCL command procedures to obtain information about the shadow set. Instead, consider using the F\$GETDVI lexical function to obtain many of the items displayed by the SHOW SHADOW command.

Furthermore, the behavior of the SET SHADOW command will also change. In addition to other new qualifiers, a new /ALL qualifier will be required if SET SHADOW is used to set characteristics in all shadow sets on a system at the same time.

Please keep these changes in mind if you are writing DCL command procedures that use these new commands.

4.26.3 Dissimilar Device Shadowing (DDS): Restriction on Adding First Dissimilar Member

V7.3-2

Dissimilar device shadowing (DDS) is a new feature introduced in OpenVMS Version 7.3-2. This note describes a temporary restriction on when it is advisable to add the first dissimilar member to a shadow set that is currently mounted in the cluster.

A dissimilar member is added to an existing shadow set by using the MOUNT command. For example, the following command adds \$1\$DGA221 to shadow set DSA22:

```
$ MOUNT DSA22:/SHADOW=($1$DGA221) VOL22
```

HP advises that a dissimilar device should be added to an existing non-DDS shadow set only from a node on which the shadow set is currently mounted and only under one of the following conditions:

- The shadow set is currently mounted on only a single node.
- The dissimilar member is added by executing the MOUNT command on the system that was the first to mount the shadow set for its current mount lifetime.

Once a shadow set has dissimilar members, it can safely be mounted or remounted on any other Version 7.3-2 system in the cluster. (Version 7.3-2 is required to access a shadow set with dissimilar members.) Once one dissimilar member has been added to a given shadow set in its current mount lifetime, dissimilar members can be added or removed without restriction as long as that shadow set remains mounted.

Failure to adhere to this restriction results in a continuous level of OpenVMS distributed locking activity between the nodes that currently have the shadow set mounted. Once started, this locking activity will not cease until the shadow set is dismounted from all nodes or from all but one node. This locking activity can interfere with your ability to mount the shadow set on additional nodes and it also consumes system resources. However, this problem does not put your on-disk data at risk.

HP expects to correct this problem in a remedial kit for Version 7.3-2.

4.26.4 Write Bitmaps and Dissimilar Device Shadowing (DDS) Caution

V7.3-2

An interaction occurs between write bitmaps and dissimilar device shadowing (DDS) when Volume Shadowing for OpenVMS is used.

DDS, a new feature in OpenVMS Version 7.3-2, allows you to construct shadow sets of disk devices that are of dissimilar sizes. (For details about DDS, refer to the *HP OpenVMS Alpha Version 7.3-2 New Features and Documentation Overview* and *HP Volume Shadowing for OpenVMS*.)

Write bitmaps keep track of application writes made to a shadow set virtual unit so that a member can be returned to that virtual unit without the overhead of a full copy. A write bitmap is created when the user issues a DISMOUNT/POLICY=MINICOPY command for a shadow set member or mounts a shadow set using the MOUNT/POLICY=MINICOPY command. When this bitmap is created, its size depends on the current size of the volume.

System Management Release Notes

4.26 Volume Shadowing for OpenVMS

When a shadow set is mounted, the logical size of the shadow set virtual unit is set to the size of the smallest member unit. When a member of the shadow set is removed, the logical size of the virtual unit is recomputed based on the sizes of the remaining members of the set. Consequently, the logical size of the virtual unit may increase.

When a write bitmap is created for a shadow set, its size is determined by the current size of the shadow set virtual unit. If the virtual unit's size subsequently increases, the bitmap will not cover the entire virtual unit. If the bitmap is then used to bring back a shadow set member with a minicopy operation, the portion of the virtual unit that is not covered by the bitmap will be copied with a full copy operation.

The following example illustrates this problem:

- Shadow set DSA1: consists of these three members:

```
$1$DGA20:  (18 GB)
$1$DGA21:  (36 GB)
$1$DGA22:  (36 GB)
```

- \$1\$DGA22: is removed from the shadow set with a minicopy bitmap using the following command:

```
$ DISMOUNT/POLICY=MINICOPY $1$DGA22:
```

The write bitmap is sized for 18 GB, the current size of the shadow set virtual unit.

- \$1\$DGA20: is removed from the shadow set. To allow the file system to utilize the entire 36 GB of the remaining member, use the following command:

```
$ SET VOLUME/SIZE DSA1
```

\$1\$DGA20 can no longer be used in this shadow set because it is smaller than the new volume size.

- \$1\$DGA22: is returned to the shadow set using this command

```
$ MOUNT/SYSTEM DSA1:/SHADOW=$1$DGA22: label
```

The logical size of DSA1: remains at 36 GB; however, the bitmap covers only the first 18 GB.

- The first 18 GB of \$1\$DGA22: are copied using the minicopy bitmap; the remaining 18 GB are copied using a full copy operation.

If the removal of a smaller shadow set member is planned, removing it before removing a larger member with a minicopy bitmap will cause a larger bitmap to be created and will avoid the performance impact of a short bitmap. (In the preceding example, you would remove \$1\$DGA20: before removing \$1\$DGA22:.)

4.26.5 KZPDC (Smart Array 5300) Restrictions

V7.3-2

Volume Shadowing for OpenVMS can be used with the KZPDC controller (Smart Array 5300) subject to the restriction that all shadow set members are formed using devices composed of fault-tolerant devices, such as the following:

- RAID 1, also known as controller-based mirroring
- RAID 5, which is striping with parity

- RAID ADG (Advanced Data Guarding), which is striping with multiple parity devices

A fault-tolerant device on the KZPDC (Smart Array 5300) controller is one that can repair data errors when the media yields errors on one of the underlying LUNs.

OpenVMS Alpha Version 7.3-2 supports shadow sets with members whose total block count varies. This new feature is known as dissimilar device shadowing (DDS). DDS allows a KZPDC device to be shadowed with a device from any supported controller.

For all prior OpenVMS versions, all devices must report the same number of total blocks for HBVS to create a multiple-member shadow set. The configuration utility sets the total number of blocks on a KZPDC or MSA1000 device to the closest match that it can make to the requested size. Because the KZPDC and the MSA1000 use the same calculation, a device created on both with the same requested size will be set to the same size. This allows HBVS to create multiple-member shadow sets.

Caution

There are cases where it will not be possible to use HBVS to create a multiple-member shadow set if a fault-tolerant device is not used. For example, a single member shadow set is formed using a device (physical disk or non-fault-tolerant device). If that device subsequently develops nonrecoverable data errors, it will not be possible to use HBVS successfully to add another member to this shadow set. Once the second member is added to the shadow set, HBVS will read the entire source device and copy it to the target device. When the data error is read from the founding or source shadow set member, HBVS will attempt to force all of the current shadow set members (the source member and the copy target) to create a “bad spot”. If this request to create a bad spot fails on either shadow set member, the shadow set will be reduced to one member.

4.26.6 SHOW DEVICE/BITMAP Command Might Produce SYSTEM-F-INTDIV Errors

V7.3-2

If you attempt to display a bitmap of a shadow set that is no longer mounted on the local node, you may get a SYSTEM-F-INTDIV error if minicopy bitmaps were being used by this shadow set.

To avoid this error, do not attempt to display bitmap information unless the shadow set is mounted. Alternatively, you can try executing the command from another node in the cluster.

This restriction will be removed in a future release of OpenVMS.

4.26.7 SHOW DEVICE/BITMAP Requires LOG_IO Privilege

V7.3-2

In OpenVMS Version 7.3-2, you must have LOG_IO privilege to execute the SHOW DEVICE/BITMAP command. If you do not have LOG_IO privilege, the command shows no bitmaps active.

This temporary restriction will be corrected with a patch in the near future.

System Management Release Notes

4.26 Volume Shadowing for OpenVMS

4.26.8 Changes in Shadow Set Merge Delay Computation

V7.3-2

During an unassisted shadow set merge operation, read I/O performance available to applications is reduced by two factors:

- The need to perform data consistency checks on all read I/Os
- Contention for I/O bandwidth by the shadow set merge operation

The shadow set merge operation employs a throttling mechanism to limit the impact of merge I/O on applications. The merge process is throttled by introducing a delay between merge I/Os when system load is detected. The logic for computing this delay has been redesigned for OpenVMS Alpha Version 7.3-2. With the new merge delay computation, the default parameter settings will result in faster merge rates for some I/O controllers, such as the HSG-80. For more information, refer to the *HP Volume Shadowing for OpenVMS* manual.

4.26.9 Deleting a Fibre Channel Device from a Controller Pair

V7.3-2

If an existing Fibre Channel device is deleted from a controller pair, the user-supplied device identifier (UDID) is also deleted. If that same UDID later is assigned to a LUN located on another controller pair, and if that device is used as a member of a host-based shadow set, it will be added to the shadow set only on the first system. All other systems with that shadow set mounted will not be able to add that device.

This problem will be fixed in a host-based Volume Shadowing remedial kit for OpenVMS Version 7.3-2. See Section 1.6 for information about how to obtain remedial kits.

4.26.10 ANALYZE/DISK/SHADOW Command Behavior

V7.3-2

When you specify the /SHADOW qualifier (new in OpenVMS Version 7.3-2) with the ANALYZE/DISK_STRUCTURE command, the entire contents of a shadow set or a specified range of blocks in a shadow set are examined for discrepancies.

If a member of the shadow set experiences connectivity problems for any reason, the ANALYZE/DISK_STRUCTURE command displays the error that it received and then returns the user to the DCL prompt.

To correct the connectivity problem and run the utility again on the same shadow set, you might need to create a temporary file on the virtual unit before reissuing the ANALYZE/DISK/SHADOW command.

Additionally, this utility may report explainable discrepancies between the shadow set members if the shadow set has not undergone a full merge since the shadow set was created. This happens if the shadow set was created using the DCL command INITIALIZE/SHADOW without the /ERASE qualifier and if the disk devices had different contents. It is important to realize that this is *not* disk corruption. The blocks that are reported as different have not been written to, but they may contain stale data; the blocks reported as inconsistent may even be allocated to a file because there may be unwritten space between the file's end-of-data location and the end of the allocated space.

You can eliminate such inconsistencies by performing a full merge. To initiate a full merge, execute the DCL command SET SHADOW/DEMAND_MERGE DSAxxx. If the devices are served by controllers that support controller-based minimerge (for example, HSJ50s), this command should be issued while the shadow set is mounted on only one node within the cluster. Otherwise, a minimerge will occur, and the discrepancy may not be resolved. When you are adding members to a single member shadow set, a full copy operation will also ensure that the disk is consistent both within and outside of the file system.

If errors are reported on an ANALYZE/DISK/SHADOW command after a full merge has been executed, they should be investigated.

Also see Section 4.26.11 for another note about ANALYZE/DISK/SHADOW command behavior.

4.26.11 ANALYZE/DISK/SHADOW Command Behavior with Dissimilar Device Shadow Sets

V7.3-2

An ANALYZE/DISK/SHADOW command may also report explainable discrepancies if a full merge has not occurred since the shadow set was logically expanded after a new member was added. The following example illustrates this problem:

- Shadow set DSA1: consists of two members:
 - \$1\$DGA20: (18 GB)
 - \$1\$DGA21: (36 GB)
- A second 36-GB member, \$1\$DGA22:, is added to the shadow set with a full copy operation.
- After the copy completes, \$1\$DGA20: is removed from the shadow set.
- At this point, if the SET VOLUME/SIZE DSA1: command is executed, the shadow set virtual unit DSA1: will increase to 36 GB. Then, ANALYZE/DISK/SHADOW will report discrepancies because only the first 18 GB of the shadow set contents were copied to \$1\$DGA22:.

The discrepancies reported by ANALYZE/DISK/SHADOW are harmless because the space in question has not yet been written by applications.

Also see Section 4.26.10 for another note about ANALYZE/DISK/SHADOW command behavior.

4.26.12 Dismount of Shadow Set Member Using /MINICOPY

V7.3

In an OpenVMS Cluster configuration, if you issue a DISMOUNT command with the /MINICOPY qualifier from a client system to dismount a shadow set member, the command might fail.

Workaround

If the first DISMOUNT command fails, repeat the command, as shown in the following example:

System Management Release Notes

4.26 Volume Shadowing for OpenVMS

```
$ SHOW DEVICE DSA5555
Device      Device      Error   Volume      Free  Trans Mnt
  Name      Status      Count   Label        Blocks Count Cnt
DSA5555:    Mounted      0   $80$DKA107:  7994646      1  18
$80$DKA107: (WILD3) ShadowSetMember  0 (member of DSA5555:)
$80$DKA302: (WILD3) ShadowSetMember  0 (member of DSA5555:)
$80$DKA303: (WILD3) ShadowSetMember  0 (member of DSA5555:)
$
$
$ DISMOUNT/POLICY=MINICOPY $80$DKA302:
%DISM-W-CANNOTDMT, $80$DKA302: cannot be dismounted
%DISM-F-SRCMEM, only source member of shadow set cannot be dismounted
$
$
$ DISMOUNT/POLICY=MINICOPY $80$DKA302:
$
```

This problem will be corrected in a future release.

Programming Release Notes

This chapter provides release notes about application and system programming on OpenVMS systems.

5.1 Backup API: Journaling Callback Events Restriction

Permanent Restriction

If an application registers a callback routine for any of the journaling events, it must register a callback routine for all the journaling callback events. The following is a list of the journaling callback events:

```
BCK_EVENT_K_JOURNAL_OPEN
BCK_EVENT_K_JOURNAL_WRITE
BCK_EVENT_K_JOURNAL_CLOSE
```

Refer to the Backup API chapter in the *OpenVMS Utility Routines Manual* for more information about registering callback routines.

5.2 C Programs: Compiling with case=sensitive Settings

Permanent Restriction

If you are compiling C programs with the setting `case=sensitive`, any `#include` files in your C program specified with the `.h` file type (lowercase `h`) will not be seen and executed. In addition, if a system `#include` file specifies another `#include` file with a `.h` file type, the second `#include` file will not be seen and an error will be generated.

To avoid this behavior, compile with `case` set to `blind`. If it is necessary to use `case=sensitive`, specify any `#include` files in your C programs either with no file type (for example, `#include <stdio>`) or with an uppercase `H` file type (for example, `#include <stdio.H>`).

Note that this does not correct the scenario where system `#include` files, such as `stdlib.h`, in turn specify `#include` files with a `.h` file type and cause an error to be generated.

5.3 C Run-Time Library

The following sections describe changes and corrections to the C Run-Time Library (RTL).

5.3.1 Undefined Symbol GETADDRINFO_COMPAT43 Fixed

V7.3-2

In some circumstances, the OpenVMS Linker flagged the `GETADDRINFO_COMPAT43` symbol as undefined. This has been fixed.

Programming Release Notes

5.3 C Run-Time Library

5.3.2 64-Bit Pointers Now Supported for Some Socket Routines

V7.3-2

The following socket routines did not have 64-bit pointer support. Now they do.

```
accept
getpeername
getsockname
inet_ntop
inet_pton
```

5.3.3 fwrite Function No Longer Writes Inappropriate Records to a File

V7.3-2

In OpenVMS Version 7.3-1, a change was made to the way the C RTL wrote fixed-length records whose size was less than the record size defined for the file. In this case, the C RTL would pad the record with zeros to create a record size equal to the maximum record size defined for the file.

Because this was a change in behavior, a feature logical should have been included to control this behavior. This feature logical has now been added. This new feature logical, DECC\$WRITE_SHORT_RECORDS, supports the previous change to the fwrite function while retaining the legacy way of writing records to a fixed-length file as the default behavior:

With DECC\$WRITE_SHORT_RECORDS enabled, short-sized records are padded with zeros to align records on record boundaries.

With DECC\$WRITE_SHORT_RECORDS disabled, the legacy behavior of writing records is implemented. This is the default behavior.

5.3.4 FILE_TYPE Macro Removed from <stdio.h> Header

V7.3-2

The <stdio.h> header file contained the nonstandard macro FILE_TYPE. This macro has been removed from the <stdio.h> header file to prevent name space pollution.

5.3.5 tm_t Structure Removed from <time.h> Header

V7.3-2

The <time.h> header file included a definition of nonstandard structure tm_t. This structure definition has been removed from the <time.h> header file to prevent name space pollution.

5.3.6 <ftw.h> Header File Now Compatible with Local-Time Function ftw

V7.3-2

Recent versions of the C RTL, along with header files supplied with Version 6.5 of the C compiler, incorrectly made a local-time version of the ftw function the default when compiling with the _VMS_V6_SOURCE macro defined.

Workaround

This behavior was incompatible with previous C RTL versions, and has been restored to default to the UTC-based `ftw` function under the same circumstances. A new macro, `__LOCAL_TIME_FTW`, has been added to control whether a UTC-based or local-time-based version of `ftw` is used. To get the local-time `ftw`, you now must compile with `/DEFINE=__LOCAL_TIME_FTW`, in addition to the macros `_VMS_V6_SOURCE` and `_DECC_V4_SOURCE`.

5.3.7 Incomplete Support for 64-bit `iconv()` Removed

V7.3-2

Incomplete support for 64-bit `iconv` was mistakenly put into Spring 2002 ECO kits and OpenVMS Version 7.3-1. This gave the illusion of 64-bit `iconv` support under some circumstances (linking with `STARLET.OLB`). This was never supported and still is not supported.

Workaround

Because 64-bit `iconv` is neither functional nor supported, `DECC$_ICONV64` has been removed.

Users compiling with `/POINTER=LONG` should now expect a compiler error message such as the following:

```
%CC-W-PTRMISMATCH, In this statement, the referenced type of the
pointer value "(char ...)0" is "long pointer to char", which is
not compatible with "short pointer to char".
```

5.3.8 `dlsym` Function No Longer Fails on Images Linked with Warnings

V7.3-2

Previously, the `dlsym` function could not activate a shareable image linked against modules that had compiler warnings. This has been fixed.

5.3.9 `dlsym` Function Now Recognizes Lowercase Names

V7.3-2

Previously, library symbols containing lowercase characters could not be loaded using the `dlsym` function. This has been fixed.

More generally, the functions that dynamically load libraries (`dlopen`, `dlsym`, `dlclose`, `dlerror`) have been enhanced to provide the following capabilities:

- Support for libraries with mixed-case symbol names
- Ability to pass a full file path to `dlopen`
- Validation of the specified library name

5.3.10 `catopen` Function Now Supports UNIX® Style Path Names

V7.3-2

Previously, the `catopen` function did not accept a UNIX style path name containing a slash (/). This has been fixed. Now, if the path contains a slash character, it is treated as a UNIX style directory name and is converted to OpenVMS format.

Programming Release Notes

5.3 C Run-Time Library

5.3.11 Memory Management Functions Are Now Reentrant

V7.3-2

Previously, the memory-mapping functions `mmap`, `munmap`, `msync`, and `mprotect` were not reentrant. Now they are.

5.3.12 Angle Brackets in File Specs Now Handled Correctly

V7.3-2

Previously, on OpenVMS Version 7.3-1 or after installing ECO kits VMS73_ACRTL-V0200 or VMS722_ACRTLV0100, directory definitions that used angle brackets may not have been processed correctly, resulting in various errors related to the file or directory paths. This problem has been fixed.

5.3.13 Race Condition in `sleep` and `usleep` Is Fixed

V7.3-2

Previously, in multithreaded applications with upcalls enabled, concurrent use of the `sleep` and `usleep` functions could cause an application to hang. This has been fixed.

5.3.14 UNIX to OpenVMS Name Translation Errors Fixed

V7.3-2

The following UNIX to OpenVMS translation errors have been fixed:

- The `decc$translate_vms` function failed to translate the file name `a.b.c`. ODS-5 support handles names only in the form returned from RMS. RMS reports the file `a.b.c` as `a^b.c`, which is translated successfully.
- The `decc$translate_vms` function incorrectly translated the file name `^%` to `?` instead of `%`.
- The `decc$to_vms` function incorrectly translated the file name `%` to the OpenVMS name `"%."` instead of `^%`.
- The `decc$translate_vms` function failed to translate the name `[crtl.][]`. The function now correctly replaces `[]` with the current directory.

5.3.15 Behavior Changed for `tzset` with Invalid TZ

V7.3-2

Previously, if the TZ environment variable contained an invalid value, the `tzset` function interpreted the invalid value as a colon (`:`), which is Greenwich Mean Time (GMT). This caused OpenVMS to interpret all file times expressed in Coordinated Universal Time (UTC) as if they were recorded in GMT.

This has been changed for OpenVMS Version 7.3-2. Now, for an invalid TZ value, `/SYS$LOCALTIME` is used.

On Tru64 UNIX systems, an invalid TZ value causes `tzset` to use the file `/etc/zoneinfo/localtime`. The OpenVMS equivalent of `/etc/zoneinfo/localtime` is the system logical name `SYS$LOCALTIME`. The OpenVMS behavior is now the same as the Tru64 UNIX behavior.

5.3.16 Attempt to Perform a pread or pwrite on an Unseekable File Is Now Properly Flagged as an Error

V7.3-2

The X/Open® specification states that an attempt to perform a pread or pwrite on a file that is not seekable should result in an error. Previously, however, the C RTL implementation of the pread and pwrite functions did not return an error for a file descriptor (FD) associated with an unseekable file. Instead, the *offset* parameter was ignored and the operation was treated like a normal read/write, so the pread or pwrite call would succeed.

The pread and pwrite functions have been modified so that when the file is not a seekable file, the error code ESPIPE is returned to indicate that the FD is associated with a pipe or FIFO and thus is not seekable.

5.3.17 Open Files Removal Now Controlled With Feature Logical

V7.3-2

In OpenVMS Version 7.3-1, the behavior of the C RTL was changed so that users could call the remove function on systems supporting the lib\$delete_file qualifier LIB\$M_FIL_IGNORE_ACCESS, and could find that files open for exclusive access by another process were unlinked. Although this is POSIX conformant behavior, it needed to be controlled by a feature logical so that existing behavior would not change unexpectedly.

OpenVMS Version 7.3-2 fixes this by adding a C RTL feature logical called DECC\$ALLOW_REMOVE_OPEN_FILES. With this feature logical enabled, the POSIX conformant behavior of the remove function is allowed. If the feature logical is not enabled, the previous behavior (disallowing the removal) occurs.

5.3.18 fstat Function Now Checks Whether Disk Access Dates Were Enabled

V7.3-2

Previously, on an ODS-5 disk with DECC\$EFS_FILE_TIMESTAMPS enabled and the access dates option turned off, the fstat function returned the actual access time for the st_atime field when it should have returned the same value as the st_mtime field.

The fstat function now checks whether disk access dates were enabled.

5.3.19 Problem With Redefinition of Structures When Including <time.h> or <signal.h>

V7.3-2

There is a potential problem with redefinition of structures in programs that directly or indirectly include either the <time.h> or <signal.h> header file. Prior to OpenVMS Version 7.3-2, the marker symbol _TIMESPEC_T_ was used to indicate that the timespec structure and timespec_t typedef had already been defined. If the <time.h> or <signal.h> header file is included *before* a definition of the timespec structure that conforms to the protocol prior to OpenVMS Version 7.3-2, the compiler issues a REDEFSTRUCT error.

To correct this problem, modify the program to use the new marker symbols __TIMESPEC and __TIMESPEC_T to indicate a prior definition of the timespec structure and timespec_t typedef, respectively.

Programming Release Notes

5.3 C Run-Time Library

Definitions Prior to Version 7.3-2:

Before OpenVMS Version 7.3-2, the <timers.h> header file included the following definitions:

```
#ifndef _TIMESPEC_T
#   define _TIMESPEC_T
   typedef struct timespec {
       unsigned long    tv_sec;        /* seconds */
       long            tv_nsec;       /* nanoseconds */
   } timespec_t;
#endif
```

In this definition, the marker symbol `_TIMESPEC_T` is used to avoid redefinition of the `timespec` structure and the `timespec_t` typedef.

Definitions in Version 7.3-2:

In OpenVMS Version 7.3-2, the <timers.h> header file is modified to include the following definition:

```
#ifndef _TIMESPEC_T
#   define _TIMESPEC_T_

#   ifndef __TIMESPEC
#       define __TIMESPEC
       struct timespec {
           unsigned long    tv_sec;    /* seconds */
           long            tv_nsec;    /* nanoseconds */
       };
#   endif

#   ifndef __TIMESPEC_T
#       define __TIMESPEC_T
       typedef struct timespec timespec_t;
#   endif
#endif
```

The modification in this example performs the following functions:

- Separates the structure definition from the typedef definition.
- Preserves the `_TIMESPEC_T_` marker symbol, which continues to indicate that both the `timespec_t` typedef and the `timespec` structure have been defined.
- Adds two new marker symbols, `__TIMESPEC_T` and `__TIMESPEC`, to indicate the definition of the `timespec_t` typedef and `timespec` structure, respectively.

In addition, the standard header files <signal.h> and <time.h> are modified to include the following structure definition:

```
#   if (!defined __TIMESPEC && !defined _TIMESPEC_T_)
#       define __TIMESPEC
       struct timespec {
           unsigned long    tv_sec;    /* seconds */
           long            tv_nsec;    /* nanoseconds */
       };
#   endif
```

5.3.20 Mixed Logical Name Search Lists Now Work

V7.3-2

Previously, using mixed (rooted and nonrooted) logical name search lists worked on OpenVMS Version 7.1-2, but could generate an error for OpenVMS Version 7.2-2 and higher.

Starting with OpenVMS Version 7.3-2, this problem has been fixed with the addition of a new feature logical. The `DECC$NO_ROOTED_SEARCH_LISTS` feature logical controls how the `decc$to_vms` function resolves search-list logicals.

When the `decc$to_vms` function evaluates a UNIX-style path string and determines the first element to be a logical name, it does the following:

- For rooted logicals or devices, the function appends ":[000000]" to the logical name.
For example, if LOG1 is a rooted logical (`$ DEFINE LOG1 [directory.]`), then `/LOG1/filename.ext` translates to `LOG1:[000000]filename.ext`.
- For nonrooted logicals, the function appends just a colon (:) to the logical name.
For example, if LOG2 is a nonrooted logical (`$ DEFINE LOG2 [directory]`), then `/LOG2/filename.ext` translates to `LOG2:filename.ext`.
- If the first element is a search-list logical, the translation proceeds by evaluating the first element in the search list, and translating the path as described previously.

The preceding three cases lead to predictable, expected results.

In cases where the first element is a search list consisting of a mixture of rooted and nonrooted logicals, translating paths as described previously can result in behavior that differs from that of older OpenVMS versions (prior to Version 7.3-1), as follows:

- Before OpenVMS Version 7.3-1, the `decc$to_vms` function appended only a colon (:) regardless of the contents of the logical. For search lists consisting of a mixture of rooted and nonrooted logicals, this resulted in certain expected behaviors.
- For OpenVMS Version 7.3-1 and later, if the first element of the mixed search list is a rooted logical, then `decc$to_vms` appends `:[000000]` to the logical name, resulting in behavior that differs from that of OpenVMS releases prior to Version 7.3-1.

`DECC$NO_ROOTED_SEARCH_LISTS` controls how the `decc$to_vms` function resolves search-list logicals and provides a means to restore the pre-Version 7.3-1 behavior.

With `DECC$NO_ROOTED_SEARCH_LISTS` enabled, the following actions occur:

- If a logical is detected in a file specification, and it is a search list, a colon (:) is appended when forming the OpenVMS file specification.
- If it is not a search list, the behavior is the same as with `DECC$NO_ROOTED_SEARCH_LISTS` disabled.

Enabling this feature logical provides the pre-Version 7.3-1 behavior for search-list logicals.

Programming Release Notes

5.3 C Run-Time Library

With `DECC$NO_ROOTED_SEARCH_LISTS` disabled, the responses are as follows:

- If a logical is detected in a file specification, and it is a rooted logical (or a search list whose first element is a rooted logical), then `:[000000]` is appended when forming the OpenVMS file specification.
- If it is a nonrooted logical (or a search list whose first element is a nonrooted logical), then just a colon (`:`) is appended.

Disabling this feature logical provides the behavior for OpenVMS Version 7.3-1 and later.

5.3.21 Program Deadlock Problem is Fixed

V7.3-2

Previously, a deadlock could occur when one process received a signal from another process, resulting in unexpected program hangs.

This problem was reported with a particular Oracle application, but could be seen by other programs as well.

This problem has been fixed.

5.3.22 `inet_ntop` function Now Defined as Returning "const char *"

V7.3-2

Previous versions of the C RTL defined the `inet_ntop` function as returning type `"char *"`. To conform to industry standards, the declaration of `inet_ntop` has been changed to return `"const char *"`.

5.3.23 `exec` No Longer Gets 0x35DF94 Error

V7.3-2

Previously, C programs that make use of `vfork/exec*` calls could see unexpected `%X35DF94` errors returned from the `exec*` call.

This problem has been fixed.

5.3.24 Compiler Errors Gone After Upgrading to OpenVMS Version 7.3-1

V7.3-2

Compiler errors no longer occur after upgrading to OpenVMS Version 7.3-1 or later.

5.3.25 Missing Entry Points for BSD 4.4 Functions Fixed

V7.3-2

Previously, the `getaddrinfo`, `freeaddrinfo`, `getnameinfo`, and `gai_strerror` functions were erroneously prefixed with `__bsd44_`, resulting in a link-time undefined-symbol error such as the following:

```
%LINK-W-USEUNDEF, undefined symbol __BSD44_GETNAMEINFO referenced
```

This problem is fixed in the `<netdb.h>` header file that ships with OpenVMS Version 7.3-2, so that these functions are no longer prefixed with `__bsd_44_`.

5.3.26 Correct IPv6 Symbol Names Now in UCX\$IPC_XHR.EXE

V7.3-2

Nine IPv6 related functions (inet6_xxxx) that were added to the C RTL for OpenVMS Version 7.3 had names that did not match the corresponding function names in the underlying TCP/IP layer. User applications that used these functions in the C RTL got an error indicating that the functions were not implemented. This problem has been fixed.

5.3.27 Translation Problem with Certain UNIX File Names Is Fixed

V7.3-2

Previously, UNIX file names such as /logname were not translated successfully when logname resolved to a file like sys\$login:login.com. This is now fixed.

5.3.28 rename Function Now Handles Logical Names Correctly

V7.3-2

In some recent versions of the C RTL, the rename function did not allow conversion to a directory specification when an ambiguous file specification was passed as a logical name for the second argument. The ambiguity was about whether the logical represented a UNIX pathname or an OpenVMS directory specification. This was a change in behavior from earlier OpenVMS versions, which did handle these rename operations correctly.

This problem has now been fixed with a new feature logical, DECC\$RENAME_ALLOW_DIR.

Enabling DECC\$RENAME_ALLOW_DIR restores the earlier OpenVMS behavior of the rename function by allowing conversion to a directory specification where the second argument is an ambiguous file specification passed as a logical name. Consider the following example with DECC\$RENAME_ALLOW_DIR enabled:

```
rename("file.ext","logical_name") /* where logical_name = dev:[dir.subdir] */  
                                  /* and :[dir.subdir] exists
```

This code results in the following construction:

```
dev:[dir.subdir]file.ext
```

This example renames a file from one directory into another directory, which is the same behavior as in versions of OpenVMS prior to Version 7.3-1. Also in this example, if dev:[dir.subdir] does not exist, rename returns an error.

Disabling DECC\$RENAME_ALLOW_DIR provides a more UNIX compliant conversion of the **logical_name** argument of rename. Consider the following example with DECC\$RENAME_ALLOW_DIR disabled:

```
rename("file.ext","logical_name") /* where logical_name = dev:[dir.subdir] */
```

This results in the following construction:

```
dev:[dir]subdir.ext
```

This example renames the file using the subdir part of the **logical_name** argument as the new file name because on UNIX systems, renaming a file to a directory is not allowed. So rename internally converts **logical_name** to a file name, and dev:[dir]subdir is the most reasonable conversion it can perform.

Programming Release Notes

5.3 C Run-Time Library

This new feature switch has a side effect of causing rename to a directory to take precedence over rename to a file. Consider this example:

```
rename ( "file1.ext","dir2" ) /* dir2 is not a logical */
```

With DECC\$RENAME_ALLOW_DIR disabled, this example results in dir2.ext, regardless of whether subdirectory [.dir2] exists.

With DECC\$RENAME_ALLOW_DIR enabled, this example results in dir2.ext only if subdirectory [.dir2] does not exist. If subdirectory [.dir2] does exist, the result is [.dir2]file1.ext.

Notes

- If DECC\$RENAME_NO_INHERIT is enabled, UNIX compliant behavior is expected, so DECC\$RENAME_ALLOW_DIR is ignored, and renaming a file to a directory is not allowed.
 - On failure, the rename function returns -1 and sets errno to one of the following values:
 - EISDIR — The new argument points to a directory, and the old argument points to a file that is not a directory.
 - EEXIST — The new argument points to a directory that already exists.
 - ENOTDIR — The old argument names a directory, and the new argument names a nondirectory file.
-

5.3.29 IPv6 Programming Note

V7.3-2

The most current TCP/IP header files ship with HP TCP/IP Services for OpenVMS Version 5.4. When Version 5.4 of TCP/IP Services is installed, the files are located in TCPIP\$EXAMPLES. The one exception is <in6.h>. In the DECC\$RTLDEF.TLB shipped with OpenVMS Version 7.3-2 <in6.h> is conceptually, but not textually, the same as the TCP/IP Services Version 5.4 variant in TCPIP\$EXAMPLES.

5.3.30 glob Function Not Supported with /POINTER_SIZE=LONG

V7.3-2

The glob function is not supported when it is compiled with the /POINTER_SIZE=LONG qualifier. Unpredictable run-time errors can occur.

5.3.31 New 64-bit Function Support: Special Considerations

V7.3-2

The following functions offer new 64-bit pointer support in OpenVMS Version 7.3-2:

getaddrinfo	getpwnam
freeaddrinfo	getpwuid
sendmsg	getpwent
recvmsg	

These functions previously offered only 32-bit support, even when compiled with `/POINTER_SIZE=LONG`. In order to preserve the previous behavior of 32-bit pointer support in these functions even when compiled with `/POINTER_SIZE=LONG`, these functions do not follow the normal convention for 32-bit and 64-bit support as documented in the *HP C Run-Time Library Reference Manual for OpenVMS Systems*.

To provide the new 64-bit support, the following variants of these functions and the corresponding structures they use have been introduced:

Function	Structure
-----	-----
__getaddrinfo32	__addrinfo32
__getaddrinfo64	__addrinfo64
__freeaddrinfo32	__addrinfo32
__freeaddrinfo64	__addrinfo64
__recvmsg32	__msghdr32
__recvmsg64	__msghdr64
__sendmsg32	__msghdr32
__sendmsg64	__msghdr64
__32_getpwnam	__passwd32
__64_getpwnam	__passwd64
__32_getpwuid	__passwd32
__64_getpwuid	__passwd64
__32_getpwent	__passwd32
__64_getpwent	__passwd64

When compiling the standard versions of these functions, you get the the following behavior:

- With `/POINTER_SIZE=32` specified, the compiler converts the call to the 32-bit version of the function. For example, `getaddrinfo` is converted to `__getaddrinfo32`.
- With `/POINTER_SIZE=64` specified, the compiler converts the call to the 64-bit version of the function. For example, `getaddrinfo` is converted to `__getaddrinfo64`.
- When the `/POINTER_SIZE` qualifier is not specified, neither the 32-bit-specific nor the 64-bit-specific function prototypes are defined.

However, a similar conversion of the corresponding structures does *not* occur for these functions. This behavior is necessary because these structures existed prior to OpenVMS Version 7.3-2 as 32-bit versions only, even when compiled with `/POINTER_SIZE=LONG`. Implicitly changing the size of the structure could result in unexpected run-time errors.

When compiling programs that use the standard version of these functions for 64-bit support, you must use the 64-bit-specific definition of the related structure. With `/POINTER_SIZE=64` specified, compiling a program with the standard function name and standard structure definition will result in compiler PTRMISMATCH warning messages.

For example, the following program uses the `getaddrinfo` and `freeaddrinfo` routines, along with the standard definition of the `addrinfo` structure. Compiling this program results in the warning messages shown in this example:

Programming Release Notes

5.3 C Run-Time Library

```
$ type test.c
#include <netdb.h>

int main ()
{
    struct addrinfo *ai;

    getaddrinfo ("althea", 0, 0, &ai);
    freeaddrinfo (ai);
    return 0;
}

$ cc /pointer_size=64 TEST.C
    getaddrinfo ("althea", 0, 0, &ai);
    ....
%CC-W-PTRMISMATCH, In this statement, the referenced type of the pointer value
"&ai" is "long pointer to struct addrinfo", which is not compatible with "long
pointer to struct __addrinfo64".
at line number 7 in file TEST.C;1
    freeaddrinfo (ai);
    ....
%CC-W-PTRMISMATCH, In this statement, the referenced type of the pointer value
"ai" is "struct addrinfo", which is not compatible with "struct __addrinfo64".
at line number 8 in file TEST.C;1
$
```

When compiling for 64 bits, you must use the 64-bit-specific version of the related structure. In the previous example, the declaration of the ai structure could be changed to the following:

```
struct __addrinfo64 *ai;
```

Or, to provide flexibility between 32-bit and 64-bit compilations, the ai structure could be declared as follows:

```
#if __INITIAL_POINTER_SIZE == 64
    struct __addrinfo64 *ai;
#else
    struct __addrinfo32 *ai;
#endif
```

5.4 Common Data Security Architecture (CDSA) Considerations

V7.3-2

Installation and Initialization

Installation of CDSA is done automatically when you install the operating system. However, you must be aware of the following considerations:

- Although the CDSA installation is part of the OpenVMS Alpha Version 7.3-2 installation, the setup and initialization of CDSA must be performed separately. Before you can use CDSA, you must perform the following manual procedure. You need the SYSPRV privilege to run this procedure. Enter the following command:

```
$ @SYS$STARTUP:CDSA$INITIALIZE
```

When a new version of CDSA is installed (for example, in an upgrade from field test to a production version, or to a new version of OpenVMS), you must run the CDSA upgrade procedure (@SYS\$STARTUP:CDSA\$UPGRADE). You should shut down any CDSA application before you run the upgrade procedure.

Programming Release Notes

5.4 Common Data Security Architecture (CDSA) Considerations

It is not necessary to rerun the initialization or upgrade procedures when the system is rebooted. You also do not need to add the initialization or upgrade procedures to the OpenVMS startup procedures.

- Do not attempt to remove CDSA from your system. Use of the PCSI command `PRODUCT REMOVE` is not supported for CDSA, even though there is an apparent option to remove CDSA. (This option is due to the use of PCSI in the installation.) CDSA is installed together with the operating system and is tightly bound with it. An attempt to remove it from OpenVMS will not work cleanly and could create other undesirable effects. An attempt to remove CDSA will result in a message similar to the following:

```
%PCSI-E-HRDREF, product CPQ AXPVMS CDSA Vn.n is referenced
by DEC AXPVMS OPENVMS V7.3-2
-PCSI-E-HRDRF1, the two products are tightly bound by this
software dependency
```

5.5 CONVERT-I-SEQ Error on CONVERT/NOSORT with Collated Key

V7.3

This potential change in behavior is restricted to a `CONVERT` command that specifies both the `/NOSORT` qualifier and a collated key type on one of the keys in the output file.

The `/NOSORT` qualifier on a `CONVERT` command indicates that the primary key is already in sorted order in the input file and directs the Convert utility not to sort it. Prior to OpenVMS Version 7.3, the Convert utility had a defect that caused it to always sort the input file if some key specified for the output file had a collated key type, regardless of whether `/NOSORT` was specified. As of OpenVMS Version 7.3, the Convert utility has been fixed to appropriately obey the `/NOSORT` qualifier on the command line, even if one of the keys in the output file is a collated key.

This means that a convert operation that previously succeeded as a side-effect of the collated key defect may now produce `%CONVERT-I-SEQ` messages if the input file is not already in sorted order by the primary key and `/NOSORT` is specified on the command line. The `/NOSORT` qualifier should not be used if the input file is not already in sorted order by the primary key.

5.6 Debugger: Previous Versions of Client/Server Interface Not Supported

V7.3

The OpenVMS Version 7.3 debugger does not support previous versions of the client/server interface. You must install the client/server interface found in the kit on the distribution media, as identified in the following table:

CPU	Operating System	Client Kit
Intel	Microsoft Windows 95, 98, NT, Me, 2000, XP	[DEBUG_CLIENTS011.KIT]DEBUGX86011.EXE
Alpha	Microsoft Windows NT	[DEBUG_CLIENTS011.KIT]DEBUGALPHA011.EXE

These client kits are self-extracting .EXE files.

Programming Release Notes

5.6 Debugger: Previous Versions of Client/Server Interface Not Supported

Once the appropriate executable file has been transferred to the PC, you can run the file to install the debug client on the PC. The InstallShield installation procedure guides you through the installation.

By default, the debug client is installed in the \Program Files\OpenVMS Debugger folder. You can also click Browse to select an alternate location.

5.7 Debugging Modes: Avoiding CPUSPINWAIT Bugchecks

V7.3-1

The OpenVMS operating system has a number of special modes of operation designed to help you debug complex hardware and software problems. In general terms, these special modes enable an extra level of tracing, data recording, and consistency checking that is useful in identifying a failing hardware or software component. These modes of operation are controlled by several system parameters: MULTIPROCESSING, POOLCHECK, BUGCHECKFATAL, and SYSTEM_CHECK.

If you are using one of these special modes (for example, to debug a device driver or other complex application), under certain conditions, generally related to high I/O loads, it is possible to incur a CPUSPINWAIT bugcheck. Specifically, any privileged code that runs for extended periods of time while holding a spinlock can cause a CPU spinwait bugcheck. Spinlocks are used to delineate the entry and exit points for critical sections, and should not be held continuously, as can occur in this situation.

To prevent a CPUSPINWAIT bugcheck, either use the system default settings for these system parameters, or reduce the loading of the system.

If you have reason to change the default settings, you can reduce the likelihood of encountering a problem by setting the SMP_LNGSPINWAIT system parameter to a value of 9000000.

5.8 HP COBOL Run-Time Library (RTL) for OpenVMS Alpha

V7.3-2

The HP COBOL RTL (DEC\$COBRTL) has been updated to V2.8-670 for OpenVMS Alpha Version 7.3-2.

5.8.1 COBOL RTL and UNSTRING

V7.3-2

The COBOL RTL now correctly handles UNSTRING OR for cases where there is overlap in the DELIMITED BY strings.

5.8.2 COBOL RTL Detects Device Full on CLOSE

V7.3-2

The COBOL RTL now detects when the device is full on a CLOSE operation.

5.8.3 COBOL RTL and Record-Locking Restriction

V7.3-1

In certain circumstances with the START or WRITE statement, COBOL acquires multiple record locks with automatic record locking. If you encounter this situation, you can clear the record locks using either UNLOCK ALL RECORDS or CLOSE followed by OPEN.

5.9 Hypersort Utility

V7.3-2

The following notes concern the Hypersort utility. Hypersort has been updated for OpenVMS Alpha Version 7.3-2. The new version of Hypersort is V04-004.

As always, continue to use SORT32 as a workaround for any unresolved problems with Hypersort or any functionality not implemented in Hypersort.

5.9.1 CONVERT Performance with Large Files

V7.3-2

Hypersort now correctly handles CONVERT with large files. Previously, just a small change in working set extent could result in very poor Hypersort performance.

5.9.2 Hypersort and VFC Files Restriction

V7.3-2

Use of VFC files with Hypersort requires /FORMAT=RECORD_SIZE:n.

5.9.3 SORT and MERGE Can Hang with Large Input Files

V7.3-2

In some cases when Hypersort is dealing with a very large input file (typically greater than 4GB), Hypersort can hang or terminate with ACCVIO. If you encounter an ACCVIO or hang with Hypersort, the workaround is to use SORT32.

5.9.4 /FORMAT=RECORD_SIZE Restriction

V7.3-1

Hypersort supports /FORMAT=RECORD_SIZE:n for use with both SORT and MERGE, with the following two restrictions:

- In all cases, if the command-specified RECORD_SIZE is less than the longest record size (LRL) of any record in the input files, the records that are too long are truncated to the RECORD_SIZE size in the sorted output file and the diagnostic message %SORT-E-BAD_LRL is issued. In this situation, the output file should be discarded and the sort should be rerun. The RECORD_SIZE parameter for the SORT command should be revised to a value appropriate to the size of the largest record as shown in the listing of a DIR/FULL command for the input files.
- SORT and MERGE produce output sequential files from input indexed files. The %SORT-E-BAD_LRL diagnostic message can also be issued for this case.

5.9.5 Using Hypersort with Search Lists and Other Uses of Logical Names

V7.3-1

Hypersort does not fully support search lists and logical names used for input files and work files. SORT32 should be used when this is a problem.

Programming Release Notes

5.9 Hypersort Utility

5.9.6 Lack of Free Space for Work Files

V7.3-1

Hypersort does not properly terminate if free space is exhausted in all available sort work files. To avoid this restriction, allocate sufficient free space for the devices used for sort work files; or use SORT32 to detect that work file space has been exhausted.

5.9.7 Hypersort and SORT32 Performance: Working Set and Page File Quota

V7.3-1

Hypersort and SORT32 use different sorting and work file algorithms. Either sort utility may be faster depending on the input file and the memory/disk/CPU configuration. Make sure that working set extent is at most one third of page file quota with either Hypersort or SORT32.

5.9.8 Hypersort and SORT32 Performance with Variable Length Records

V7.3-1

SORT32 and Hypersort allocate fixed sized slots for sort work files based on the longest record length (LRL) information in the file. To improve sort performance, try to set LRL information in the file as close as possible to the actual longest record length. Poor initial performance may be the result of sorting some files produced by C programs, because the LRL is set higher than needed (to 32767).

5.9.9 Input Asterisk (*) Restriction

V7.3

Hypersort does not support asterisk (*) as an input file specification.

5.9.10 Work File Directories Restriction

V7.3

Hypersort work files must be redirected to directories that allow multiple file versions that can handle the number of requested work files. This restriction also exists in SORT32.

5.10 Librarian Utility: PGFLQUOTA Should Exceed 23000

V1.5

The OpenVMS Alpha LIBRARIAN sometimes does not inform you of errors during compression, data reduction, or data expansion operations. This problem occurs if the account or process in which the LIBRARIAN is running has a low PGFLQUOTA process quota. Operation failure is not readily apparent because the \$PUTMSG system service always returns a status of SS\$_NORMAL, even when the system service fails. However, when a failure occurs, the LIBRARIAN returns a status other than Success.

To work around this problem, run the compression, data reduction, or data expansion operation in an account with a PGFLQUOTA process quota greater than 23000. In addition, ensure that your command procedures check the return status from the LIBRARY command.

5.11 Lightweight Directory Access Protocol (LDAP) API

The following sections contain release notes pertaining to the LDAP API.

5.11.1 The Routine `ldap_get_option` Returns Error -1 When `ld` Is NULL

V7.3

Using a value of NULL for the `ld` parameter in a call to `ldap_get_options()` results in an error of -1 being returned, rather than the routine returning a set of global default data.

5.11.2 The Routine `ber_flatten()` Fails to Detect Mismatched Braces

V7.3

The routine `ber_flatten()` does not correctly detect the situation where `'{'` and `'}'` format modifiers in a `BerElement` are incorrectly matched.

5.12 Linker Utility

The following sections describe release notes pertaining to the Linker Utility.

5.12.1 Linker Appears to Hang When Many Files Are Specified

V7.3-2

When the `RMS_RELATED_CONTEXT` linker option is on (the default is `RMS_RELATED_CONTEXT=YES`) and a nonexistent file is specified in a list of files for the `LINK` command, the linker's call to `LIB$FIND_FILE` takes a long time to complete and the linker may appear to hang. Depending on the number of files being linked and the use of logical names in their specification, the linker may take hours to finish because `LIB$FIND_FILE` locates every combination of the missing file's prefix before displaying a "file not found" message. Note that you cannot terminate the linker process with `Ctrl/Y` after the linker has called `LIB$FIND_FILE`.

Workaround

Specify `/OPTION` in the `LINK` command. When you press `Return`, the linker waits for you to enter information into an options file. When you are finished, press `Ctrl/Z`. To avoid the problem described in this release note, include the following items in the options file:

- On the first line, specify this option:

```
RMS_RELATED_CONTEXT=NO
```

With the `RMS_RELATED_CONTEXT` option set to `NO`, any missing file listed in this options file will generate an immediate "file not found" message.

- On subsequent lines, specify the files to be linked, using full file specifications (in the form `disk:[dir]filename.ext`) for every file. Full file specifications are required because when you specify `RMS_RELATED_CONTEXT=NO`, file name "stickiness" is disabled.

For example, consider the following `LINK` command:

```
$ LINK DSK:[TEST]A.OBJ, B.OBJ
```

Programming Release Notes

5.12 Linker Utility

If you want to specify this command with RMS_RELATED_CONTEXT=NO, you would specify /OPTION and then enter full file specifications for the files to be linked, as follows:

```
$ LINK SYS$INPUT:/OPTION
RMS_RELATED_CONTEXT=NO
DSK:[TEST]A.OBJ, DSK:[TEST]B.OBJ Ctrl/Z
$
```

For more information about the RMS_RELATED_CONTEXT option, refer to the *OpenVMS Linker Utility Manual*.

Example

The following example shows how the linker appears to hang when file DOES_NOT_EXIST.OBJ is included in the list and the RMS_RELATED_CONTEXT option is not specified (and therefore defaults to YES).

```
$ DEFINE DSK$ WORK4:[TEST.LINKER.OBJ.]
$ DEFINE RESD$ ROOT$, ROOT2$, ROOT3$, ROOT4$, ROOT5$, DISK_READ$:[SYS.] ❶
$ DEFINE ROOT$ WORK4:[TEST.PUBLIC.TEST]
$ DEFINE ROOT2$ WORK4:[TEST.LINKER.]
$ DEFINE ROOT3$ WORK4:[TEST.UTIL32.]
$ DEFINE ROOT4$ WORK4:[TEST.PUBLIC.]
$ DEFINE ROOT5$ WORK4:[TEST.PUBLIC.TMP]
$ LINK/MAP/FULL/CROSS/EXE=ALPHA.EXE RESD$:[TMP OBJ] A.OBJ, -
_$ RESD$:[SRC]B.OBJ,C,DSKD$:[OBJ]D.OBJ,E,RESD$:[TMP SRC]F.OBJ, -
_$ RESD$:[TEST]G.OBJ,RESD$:[SRC.OBJ]H,RESD$:[COM]DOES_NOT_EXIST.OBJ
Ctrl/T NODE6::_FTA183: 15:49:46 LINK CPU=00:02:30.04 PF=5154 IO=254510 MEM=134 ❷
Ctrl/T NODE6::_FTA183: 15:49:46 LINK CPU=00:02:30.05 PF=5154 IO=254513 MEM=134
Ctrl/T NODE6::_FTA183: 15:50:02 LINK CPU=00:02:38.27 PF=5154 IO=268246 MEM=134
Ctrl/T NODE6::_FTA183: 15:50:02 LINK CPU=00:02:38.28 PF=5154 IO=268253 MEM=134
Ctrl/T NODE6::_FTA183: 15:50:14 LINK CPU=00:02:44.70 PF=5154 IO=278883 MEM=134
```

- ❶ This command defines logical names and equivalents.
- ❷ Each time you press Ctrl/T, the CPU and IO values increase, but the MEM and PF values do not, indicating that LIB\$FIND_FILE is being called.

As shown in the following example, using an options file to set RMS_RELATED_CONTEXT to NO causes the link operation to finish immediately when it encounters the missing file.

```
$ DEFINE DSK$ WORK4:[TEST.LINKER.OBJ.]
$ DEFINE RESD$ ROOT$, ROOT2$, ROOT3$, ROOT4$, ROOT5$, DISK_READ$:[SYS.]
$ DEFINE ROOT$ WORK4:[TEST.PUBLIC.TEST.]
$ DEFINE ROOT2$ WORK4:[TEST.LINKER.]
$ DEFINE ROOT3$ WORK4:[TEST.UTIL32.]
$ DEFINE ROOT4$ WORK4:[TEST.PUBLIC.]
$ DEFINE ROOT5$ WORK4:[TEST.PUBLIC.TMP.]
$ LINK/MAP/FULL/ CROSS /EXE=ALPHA.EXE SYS$INPUT:/OPTION
RMS_RELATED_CONTEXT=NO
RESD$:[TMP OBJ]A.OBJ,RESD$:[SRC]B.OBJ,RESD$:[SRC]C,DSKD$:[OBJ]D.OBJ
DSKD$:[OBJ]E,RESD$:[TMP SRC]F.OBJ,RESD$:[TEST]G.OBJ
RESD$:[SRC.OBJ]H,RESD$:[COM]DOES_NOT_EXIST.OBJ Ctrl/Z

%LINK-F-OPENIN, error opening DISK_READ$:[SYS.] [COM]DOES_NOT_EXIST.OBJ; as input
-RMS-E-FNF, file not found
$
```

5.12.2 Change in Linker Default Behavior with Library Check

V7.3-1

Previously, the linker's check between the library and the shareable image was too sensitive. It compared against the exact date and time, signaling LINK-I-DATMISMCH, if no match was found. Now, however, it makes only the same check that the image activator does: that is, it uses the GSMATCH criteria to verify compatibility.

The old behavior (check for date and time) can be obtained by setting the logical name LINK\$SHR_DATE_CHECK.

5.12.3 Limit of 25 Elements on Stack

Permanent Restriction

Developers who are creating object files should be aware that the linker's internal stack is guaranteed for only 25 elements. Any calculations must be done within this constraint.

5.13 LTDRIVER: CANCEL SELECTIVE Restriction

Permanent Restriction

In releases prior to OpenVMS Version 6.1, LTDRIVER did not set the "extended DDT" bit; therefore, the POSIX function CANCEL SELECTIVE did not work with LTDRIVER. This problem has been corrected, but a restriction remains.

Although this fix allows \$QIO reads and writes to be selectively canceled, any \$QIO done to the port driver (that is, with the IO\$_TTY_PORT function modifier — such as a LAT connect \$QIO) *cannot* be canceled with CANCEL SELECTIVE.

5.14 Mail Utility: Threads Restriction for Callable Mail

V7.1

OpenVMS callable mail routines are *not* thread-safe. Refer to the *Guide to the POSIX Threads Library* for more information about calling non-thread-safe routines within a threaded application.

Because callable mail context information is maintained on a per-process (rather than a per-thread) basis, multiple threads performing context-based processing must be synchronized so that only *one* mail context of a given type is active at once. Otherwise, one thread could corrupt another thread's mail operations.

On OpenVMS Alpha systems, there is an additional restriction when kernel threads is enabled in a multithreaded environment. In this environment, callable mail should be used only in the initial thread.

5.15 POSIX Threads Library

The following sections contain release notes pertaining to the POSIX Threads Library (formerly named DECthreads).

Also see Section A.3 for a related note.

Programming Release Notes

5.15 POSIX Threads Library

5.15.1 C Language Compilation Header File Changes

V7.3-2

The following changes have been made in OpenVMS Version 7.3-2:

- INTS.H

In some prior releases of OpenVMS, the POSIX Threads C language header file `PTHREAD_EXCEPTION.H` inadvertently contained a `#include` of the C header file `INTS.H`. This has been corrected in OpenVMS Version 7.3-2. `PTHREAD_EXCEPTION.H` no longer causes `INTS.H` to be included in a compilation. There may be some applications whose compilation requires the presence of `INTS.H` and which are erroneously relying on `PTHREAD_EXCEPTION.H` to provide it.

Recompiling such application source modules on an OpenVMS Version 7.3-2 system will result in diagnostic messages from the C compiler. These messages identify symbols or data types (for example, `int32`) that originate in `INTS.H` and are undefined. To correct such application source modules, add a direct `#include` of `<ints.h>` before the first use of the corresponding symbols or types.

- `timespec_t` typedef

In prior releases of OpenVMS, the POSIX Threads C language header file `PTHREAD.H` contained a definition for a typedef named `timespec_t`. This is a nonstandard symbol, which does not belong in `PTHREAD.H`. (This typedef was present for historic reasons related to the contents of C RTL header files such as `TIME.H` and `TIMERS.H`.) For OpenVMS Version 7.3-2, the standards compliance of the C RTL and threads header files has been improved. As a result, `PTHREAD.H` no longer provides the `timespec_t` typedef.

There may be some applications whose compilations require the `timespec_t` typedef, and which erroneously rely on `PTHREAD.H` to provide it—either directly or indirectly (for example, by using `TIS.H`). If such an application source module is recompiled on an OpenVMS Version 7.3-2 system, you may get C compiler diagnostic messages listing `timespec_t` as an unknown symbol or type. To correct such application source modules, either replace the uses of `timespec_t` with structure `timespec`, or include the C RTL header file `TIMERS.H` before the first use of the `timespec_t` symbol.

If your application build environment uses a private copy of any older C RTL or threads header files or an extract of them that includes the `timespec` structure or the `timespec_t` typedef (both of which are not recommended), you may see an additional compilation error. The compiler may display messages stating that the `timespec` structure is redefined or defined twice. In such a case, revert to using the system-supplied C RTL and threads header files, or replace the private extracts involving the `timespec` structure with an inclusion of the system-supplied `TIME.H` header file.

5.15.2 New Priority Adjustment Algorithm

V7.3-2

As of OpenVMS Version 7.3-2, the adaptive thread scheduling behavior that is described in the *Guide to the POSIX Threads Library* has been implemented with a new priority adjustment algorithm. In some cases, the new algorithm should help avoid problems that can arise when throughput-policy threads of different priorities share synchronization objects. Priority adjustment can also improve

application throughput and overall system utilization. Priority adjustment of threads with throughput scheduling policy is automatic and transparent.

5.15.3 Process Dumps

V7.3

If the POSIX Threads Library detects an uncorrectable serious problem at run time (such as data structures that have been damaged by data corruption somewhere in the application), the library may terminate the running image. During termination, the library may trigger creation of a process dump file (which can subsequently be used to diagnose the failure, by way of ANALYZE/PROCESS_DUMP). The size of such a process dump file depends on the size of the process's address space at the time of the failure and can be quite large.

5.15.4 Dynamic CPU Configuration Changes

V7.3

Starting in OpenVMS Version 7.3, the POSIX Threads Library is sensitive to dynamic changes in the number of CPUs that are configured for a running multiprocessor Alpha system. When use of multiple kernel threads is enabled (by way of the LINK/THREADS_ENABLE qualifier or the THREADCP command verb) for an image, the POSIX Threads Library monitors the apparent parallelism of an application and creates multiple kernel threads up to the number of CPUs available. Each kernel thread can be scheduled by the OpenVMS executive to execute on a separate CPU and, therefore, can execute simultaneously.

While an application is running, an operator can stop or start a CPU. Such a dynamic change affects the allowable number of kernel threads that future image activations can create. It also will now affect images that are currently executing.

When a CPU is added or removed, the threads library will query for the new number of active CPUs, and compare this to the number of kernel threads that the process is currently using. If there are now more CPUs than kernel threads, the library will try to spread out the existing POSIX threads over the CPUs (creating new kernel threads as needed, now or in the future). If there are now fewer CPUs than kernel threads, the library will force the extra kernel threads to hibernate, and will reschedule the POSIX threads onto the remaining kernel threads. This will ensure that — so far as the process is concerned — there will not be more kernel threads competing for CPU resources than are available.

5.15.5 Enhanced Debugging of Threaded Programs

V7.3

The POSIX Threads Library provides enhanced data collection capabilities to support monitoring and debugging tools. These capabilities provide support for Visual Threads, a new debugging and analysis tool for threaded programs on OpenVMS Alpha systems. Visual Threads, which is licensed with OpenVMS Version 7.3, provides monitoring, automatic debugging, and performance evaluation of multithreaded applications.

Programming Release Notes

5.15 POSIX Threads Library

5.15.6 Debugger Metering Function Does Not Work

V7.0

The metering capability of the POSIX Threads debugger does not work.

If you use the procedure to debug a running program that is described in Section C.1.1 of the *Guide to the POSIX Threads Library*, your process could fail with an ACCVIO message.

5.16 Privileged Interfaces and Data Structures

This section contains release notes concerning privileged code and data structures.

5.16.1 Per-Thread Security Impacts Privileged Code and Device Drivers

V7.3-1

The method used for attaching a security profile to an I/O Request Packet (IRP) changed with Version 7.2.

In versions of OpenVMS prior to Version 7.2, the IRP structure contained the address of the processwide Access Rights Block (ARB) security structure of the requestor. Beginning with OpenVMS Alpha Version 7.2, the address of the new security profile structure (Persona Security Block, or PSB) was added to the IRP as a functional replacement of the ARB address.

The I/O subsystem maintains its access to the PSB through a reference counter within the PSB. The I/O subsystem increments this reference counter at the time of IRP creation and decrements the counter at I/O postprocessing of that IRP. When this counter reaches zero, the PSB structure is deallocated.

Device drivers that create or clone copies of IRPs to facilitate multiple I/O operations per request, and subsequently pass the copies to the I/O subsystem for postprocessing, must make code changes to account for the extra references to the PSB in these additional IRPs. This is done by passing the PSB address located in the copied IRP to the NSA_STD\$REFERENCE_PSB routine. The include file and routine call for NSA_STD\$REFERENCE_PSB is as follows:

```
#include <security-macros.h>
/* Increment REFCNT of PSB that is now shared by both IRPs */
nsa_std$reference_psb( irp->irp$psb );
```

Device drivers need to make this change under the following conditions:

- If a device driver creates a new IRP by duplicating an existing IRP and submits both the original and the duplicate IRPs for I/O postprocessing by calling IOC_STD\$SIMREQCOM or IOC_STD\$DIRPOST1, the device driver must call NSA_STD\$REFERENCE_PSB sometime after duplicating the IRP, but before submitting it for I/O postprocessing.
- If a device driver creates a new IRP by duplicating an existing IRP and does not put the address of some procedure descriptor into the IRP\$L_PID cell in either the copy or the original IRP, and the device driver submits both the original and the duplicate IRPs for I/O postprocessing by calling IOC_STD\$REQCOM, COM_STD\$POST, COM_STD\$POST_NOCNT, or IOC_STD\$POST_IRP, the device driver must call NSA_STD\$REFERENCE_PSB sometime after duplicating the IRP but before submitting it for I/O postprocessing.

Device drivers that perform these steps are also likely to put the address of some procedure descriptor into `IRP$L_PID`. Therefore, most device drivers that duplicate IRPs should be able to function correctly on OpenVMS Version 7.2 or higher without making source changes, relinking, or recompiling.

Failure to call `NSA_STD$REFERENCE_PSB` in these circumstances will result in corrupt tracking information within the PSB, which can result in system failures.

If you make code changes in a device driver to call `NSA_STD$REFERENCE_PSB`, you must recompile and relink the driver to run on OpenVMS Version 7.2 or higher.

5.16.2 IPL Requirement For OpenVMS Fork Thread Creation Now Enforced

V7.3-1

Several routines are used by privileged code to create OpenVMS fork execution threads. These routines run in system context independent of any process. There are four variations of these routines, depending on whether an immediate or queued fork is required and on which language interface is being used:

- `EXE$QUEUE_FORK`
- `EXE_STD$QUEUE_FORK`
- `EXE$PRIMITIVE_FORK`
- `EXE_STD$PRIMITIVE_FORK`

These routines must be called at or above `IPL$_RESCHED`, to prevent accidental rescheduling to a different CPU during their execution. Such a reschedule could cause the system to hang.

In OpenVMS V7.3-1, if `SYSTEM_CHECK` is set to 1, these routines check the system IPL at entry. If the IPL is below `IPL$_RESCHED`, the system will fail with an `SPLINVIPL` bugcheck.

For performance reasons, the IPL is not verified if `SYSTEM_CHECK` is set to zero (the default). Incorrect code may cause the system to hang if a reschedule to another CPU occurs during execution of these routines from process context (for example, below `IPL$_RESCHED`).

5.17 RTL Library (LIB\$)

The following notes pertain to `LIB$FIND_IMAGE_SYMBOL`.

5.17.1 LIB\$FIND_IMAGE_SYMBOL: Error in the *OpenVMS RTL Library (LIB\$) Manual*

V7.3-1

In the *OpenVMS RTL Library (LIB\$) Manual*, there is an error in the description of the `flags` argument to `LIB$FIND_IMAGE_SYMBOL`. `Flags` is documented as passed by reference. This is incorrect and returns an error message, `LIB-F-INVARG`, as a return value. If `flags` is passed by value, `LIB$FIND_IMAGE_SYMBOL` works as expected.

This error will be corrected in the next revision of the *OpenVMS RTL Library (LIB\$) Manual*.

Programming Release Notes

5.17 RTL Library (LIB\$)

5.17.2 LIB\$FIND_IMAGE_SYMBOL Signals Warning

V7.1

LIB\$FIND_IMAGE_SYMBOL may signal a warning (LIB\$_EOMWARN) to indicate that the image being activated contains modules that had compilation warnings. A condition handler used with LIB\$FIND_IMAGE_SYMBOL should probably handle this as a special case.

To allow LIB\$FIND_IMAGE_SYMBOL to continue execution after signaling LIB\$_EOMWARN, the condition handler should exit with \$\$\$_CONTINUE. For this reason, you may choose not to use LIB\$SIG_TO_RET as a condition handler for LIB\$FIND_IMAGE_SYMBOL.

5.18 Screen Management (SMG\$) Documentation

The following information should be added to topics in the reference section at the end of the *OpenVMS RTL Screen Management (SMG\$) Manual*:

V7.2

- The following statement should be added to the Condition Values Returned section of routine SMG\$DELETE_VIRTUAL_DISPLAY:

"Any condition value returned by the \$DELPRC system service."

- The description of routine SMG\$GET_TERM_DATA contains an error in the Arguments section for the **capability-data** argument. The correction is as follows:

access: write-only
mechanism: by reference, array reference

- The description of routine SMG\$SET_OUT_OF_BAND_ASTS contains an error in the Arguments section for the **AST-argument** argument. The symbolic names in the Data Structure diagram are incorrect. The symbolic names in the paragraph under this diagram are correct. The correct and incorrect symbolic names are as follows:

Incorrect	Correct
SMG\$L_PASTEBOARD_ID	SMG\$L_PBD_ID
SMG\$L_ARG	SMG\$L_USER_ARG
SMG\$B_CHARACTER	SMG\$B_CHAR

V7.1

- In the documentation for the SMG\$READ_COMPOSED_LINE routine, the following text should be appended to the description of the **flags** argument:
"The terminal characteristic /LINE_EDITING should be set for your terminal for these flags to work as expected. /LINE_EDITING is the default."
- The description of routine SMG\$SET_KEYPAD_MODE should contain this note:

Note

Changing the keypad mode changes the physical terminal setting. This is a global change for all virtual keyboards, not just the virtual keyboard

specified by the **keyboard-id** argument.

5.19 SORT32 Utility

The following release notes concern the SORT32 utility. Note that SORT32 is recommended as a workaround for unresolved problems with Hypersort and for functionality not implemented in Hypersort.

SORT32 has been updated to V07-006 for OpenVMS Alpha Version 7.3-2.

5.19.1 VFC Files Now Written Correctly

V7.3-2

The version of SORT32 that shipped with OpenVMS Alpha Versions 7.3 and 7.3-1 did not properly write VFC information to VFC output files when specification files were used. This problem has been corrected.

5.19.2 SORT32 May Not Delete Temporary Work Files

V7.3-2

SORT32 does not always delete temporary work files. It's a good idea to periodically check SYS\$SCRATCH or wherever you put SORT32 work files to see if any undeleted work files can be deleted to free up disk space.

5.19.3 SORT/SPECIFICATION and Compound Conditions Restriction

V7.3-1

SORT32 does not issue a diagnostic for a compound condition in a key specification file not enclosed in parentheses, such as the following:

```
/Condition=(Name=Test1,  
            TEST=(Field2 EQ "X") AND (Field3 EQ "A"))
```

That condition should instead be specified as:

```
/Condition=(Name=Test1,  
            TEST=((Field2 EQ "X") AND (Field3 EQ "A")))
```

5.19.4 SORT32 and Hypersort Performance: Working Set and Page File Quota

V7.3-1

SORT32 and Hypersort use different sorting and work file algorithms. Either sort utility may be faster depending on the input file and the memory/disk/CPU configuration. Make sure that working set extent is at most one third of page file quota with either SORT32 or Hypersort.

5.19.5 SORT32 and Hypersort Performance with Variable Length Records

V7.3-1

SORT32 and Hypersort allocate fixed-sized slots for sort work files based on the longest record length (LRL) information in the file. To improve sort performance, try to set LRL information in the file as close as possible to the actual longest record length. Poor initial performance may be the result of sorting some files produced by C programs, because the LRL is set higher than needed (to 32767).

Programming Release Notes

5.19 SORT32 Utility

5.19.6 SORT32 Work File Directories Restriction

V7.3

SORT32 work files must be redirected to directories that allow multiple file versions that can handle the number of requested work files. This restriction also exists in Hypersort.

5.20 System Services: SYS\$ACM(W) Changes

V7.3-2

The functional changes to SYS\$ACM[W] described in this section are introduced in OpenVMS Version 7.3-2. In these descriptions, **nonprivileged processes** refer to processes running in user mode that do not have SECURITY privilege.

- Timeout processing
Timeout processing is now enforced for nonprivileged processes. Other processes can request timeout processing by specifying the ACME\$_TIMEOUT function modifier.
- Dialogue mode iteration limit
Nonprivileged processes are now limited in the number of iterative requests they can make in a dialogue sequence of calls.
- Logon type restriction
The following ACME\$_LOGON_TYPE item code values are reserved for use by LOGINOUT:

ACME\$_K_BATCH
Zero (0)

You also no longer need IMPERSONATE privilege to specify the ACME\$_LOGON_TYPE item code.

Refer to the SYS\$ACM[W] system service description in *HP OpenVMS System Services Reference Manual* for more details.

5.21 Timer Queue Entries (TQEs)

V7.3-1

Management of Timer Queue Entries was redesigned for OpenVMS Alpha Version 7.3-1 to provide significantly higher performance for systems using many TQEs. This change is transparent to nonprivileged applications.

Privileged code can no longer manipulate TQEs directly in any form. In particular, directly accessing pointers in the TQE's queue header (TQE\$_TQFL/TQE\$_TQBL) causes an access violation in almost all cases. Privileged code may continue to use the internal routines exe_std\$instimq/exe\$instimq and exe_std\$rmvtimq/exe\$rmvtimq to enter or remove Timer Queue Entries.

Hardware Release Notes

This chapter contains information about the following hardware products:

- ALPHAbook 1 (Section 6.1)
- AlphaServer 1000A (Section 6.2)
- AlphaServer 2100 (Section 6.3)
- AlphaServer 8200/8400 (Section 6.4)
- AlphaServer ES47/ES80/GS1280 Systems (Section 6.5)
- AlphaServer GS Series (Section 6.6)
- AlphaStation 200/400 (Section 6.7)
- AlphaStation 255 (Section 6.8)
- ATI RADEON 7500 Graphics (Section 6.9)
- DECwindows X11 Display Server (Section 6.10)
- DIGITAL Modular Computing Components (Section 6.11)
- Digital Personal Workstation (Section 6.12)
- Dual-Controller HSGnn device (Section 6.13)
- PowerStorm 300/350 PCI Graphics Controller (Section 6.14)
- RFnn DSSI disk devices (Section 6.15)
- RZnn disk devices (Section 6.16)
- ZLX graphics boards (Section 6.17)

A few notes about using device drivers are also included at the end of this appendix.

6.1 ALPHAbook 1

V7.3-2

Starting with OpenVMS Alpha Version 7.3-2, the ALPHAbook 1 notebook computer is no longer supported.

6.2 AlphaServer 1000A

The following sections contain release notes pertaining to the AlphaServer 1000A computer.

Hardware Release Notes

6.2 AlphaServer 1000A

6.2.1 Bus Probe Algorithm Default

V7.1

You cannot set the console variable `BUS_PROBE_ALGORITHM` to *OLD* on AlphaServer 1000A computers. The default setting is *NEW*. If you reset the bus probe algorithm to *OLD*, your OpenVMS system will not boot correctly.

6.2.2 Installation Failure with DEFPA Adapter

V7.1

When you attempt to install the OpenVMS operating system on an AlphaServer 1000A computer that uses a DEFPA adapter, the installation may fail, resulting in a `KERNEL STACK NOT VALID HALT` error message. If this failure occurs, powercycle your system and restart the installation.

6.3 AlphaServer 2100

The following sections contain information specific to the AlphaServer 2100 series computer.

6.3.1 Console Display

V7.2

On AlphaServer 2100 and 2100A systems, a console display similar to the following is normal and does not represent system errors:

```
P00>>>SET CONSOLE SERIAL
P00>>>INIT

VMS PALcode X5.48-112, OSF PALcode X1.35-81

starting console on CPU 0
initialized idle PCB
initializing semaphores
initializing heap
initial heap 1c0c0
memory low limit = 132000
heap = 1c0c0, 13fc0
.
.
.

probing hose 0, PCI
probing PCI-to-EISA bridge, bus 1
probing PCI-to-PCI bridge, bus 2
*** unable to assign PCI base address
*** bus 2, slot 7, function 0, size 00001000 (16 bit I/O)
bus 1, slot 1 -- fra -- DEFPA
bus 1, slot 2 -- vga -- Compaq Qvision
bus 1, slot 3 -- pua -- KFESA
bus 2, slot 1 -- pka -- NCR 53C810
bus 2, slot 6 -- pkb -- NCR 53C810
bus 2, slot 7 -- pkc -- DEC KZPSA
bus 0, slot 7 -- ewa -- DECchip 21041-AA
initializing keyboard

Memory Testing and Configuration Status
Module   Size   Base Addr   Intlv Mode   Intlv Unit   Status
-----
0        64MB   00000000    1-Way        0           Passed
Total Bad Pages 0
```



```
Testing the System
Testing the Disks (read only)
Testing the Network
econfig:          20041 99
econfig:          20042 04
econfig:          20043 00
AlphaServer 2100A Console V4.3-130, built on Oct 26 1996 at 19:44:57
P00>>>P
```

Note that in the previous display, the KZPSA adapter is successfully installed despite the error message displayed in the following lines:

```
*** unable to assign PCI base address
*** bus 2, slot 7, function 0, size 00001000 (16 bit I/O)
```

6.3.2 SCSI Controller Restriction

V6.2

The Adaptec 1740/1742 SCSI controller (PB2HA-SA) is not supported on AlphaServer 2100 systems having more than 1 gigabyte (GB) of memory. If the controller is connected to such a system, the following message appears on the operator's console:

```
%PKJDRVR-E- The direct DMA window does not map all of memory. Port is going OFFLINE.
```

6.4 AlphaServer 8200/8400: FRU Table Error

V7.2

The error log buffer size is controlled by the system parameter ERLBUFFERPAGES, which has a maximum value of 32 pagelets. If the Field Replaceable Unit (FRU) table exceeds this limit during a boot of the OpenVMS Alpha operating system on an AlphaServer 8200/8400 or 4100 system, an entry will not be written to the error log file.

6.5 AlphaServer ES47/ES80/GS1280 Systems

This section contains release notes of interest to users of AlphaServer ES47/ES80/GS1280 systems.

6.5.1 Firmware Version 6.6 Recommended

V7.3-2

OpenVMS has set the recommended firmware level for AlphaServer ES47/ES80/GS1280 systems to Version 6.6. Until that firmware is installed, you will see the following informational message when you boot:

```
%SYSBOOT-W-FIRMREV, Firmware rev. 6.5 is below the recommended minimum of 6.6.
Please update your firmware to the recommended revision,
which can be found on the firmware CD labeled:
Alpha systems firmware update 6.5.
```

You can obtain Version 6.6 of the firmware from the following AlphaServer website (URL is case sensitive):

```
ftp://ftp.digital.com/pub/Digital/Alpha/firmware/index.html
```

Hardware Release Notes

6.5 AlphaServer ES47/ES80/GS1280 Systems

6.5.2 Setting SYSGEN Parameter PHYSICAL_MEMORY

V7.3-2

Because of hardware configuration requirements on the AlphaServer ES47/ES80/GS1280 systems, HP does not recommend altering the setting of the system parameter PHYSICAL_MEMORY from its default setting of -1. Artificially reducing the amount of memory can produce unpredictable results on these systems.

6.5.3 RAD Support

V7.3-2

OpenVMS support for resource affinity domains (RADs), also known as NUMA support or awareness, has not been qualified in OpenVMS Alpha Version 7.3-2 for the AlphaServer ES47/ES80/GS1280 systems. For more information about RAD support, see the *HP OpenVMS Alpha Partitioning and Galaxy Guide*.

6.5.4 License Requirements

V7.3-2

AlphaServer ES47/ES80/GS1280 systems require a minimum of two OpenVMS software licenses: one license for base support and one license for dual SMP support for the first two processors. This is a change from the previous way of licensing OpenVMS AlphaServer SMP systems. The dual SMP licenses for OpenVMS are included with the CPU modules when you purchase an OpenVMS system or when you purchase additional CPU modules for an OpenVMS system.

6.5.5 TKNOVF Error on Output for Some F\$GETSYI Item Codes

V7.3-2

The output for some F\$GETSYI item codes on an AlphaServer ES47/ES80/GS1280 system may exceed the DCL maximum string length, resulting in a TKNOVF warning from DCL. Item codes that exhibit this behavior include GLX_MBR_NAME and GLX_MBR_MEMBER.

The problem is being investigated for possible resolution in a future release.

6.5.6 STOP/CPU and Shutdown Behavior

V7.3-2

Because of hardware restrictions, any CPU on an AlphaServer ES47/ES80/GS1280 system with an attached I/O drawer cannot be stopped by using the DCL command STOP/CPU. In contrast, CPUs on these systems without an attached I/O drawer can be stopped with this command.

When the shutdown procedure is invoked on an ES47/ES80/GS1280 system with an attached I/O drawer, an error message such as the following might be displayed:

```
%SYSTEM-W-WRONGSTATE, CPU 5 is in the wrong state for the requested operation
```

You can ignore such messages. The shutdown will complete correctly.

6.5.7 Setting Time at MBM

V7.3-2

You must set the correct time and date on the MBM of an AlphaServer ES47/ES80/GS1280 system. If you do not, OpenVMS might display an incorrect time and date.

6.5.8 ERLBUFFERPAGES Must Be Increased for AlphaServer GS1280 Systems

V7.3-2

Some AlphaServer systems capture error state when the operating system is unable to do so. When such data exists, the operating system logs the data to the error log during the next boot.

AlphaServer GS1280 systems capture a large amount of data — too large for the default error log buffer to hold. HP recommends that you set the ERLBUFFERPAGES system parameter on these systems to its maximum of 128. The setting of 128 makes each error log buffer 64K bytes. You might want to reduce the number of buffers by decreasing the ERRORLOGBUFFERS system parameter.

6.6 AlphaServer GS Series Systems

This section contains release notes of general interest to most users of the AlphaServer GS Series systems. See a related note in Section 4.22.5.

6.6.1 AlphaServer GS80/160/320 Systems: Device Restriction

Permanent Restriction

Only one set of the following devices found on the legacy bus adapter is configured and supported per partition in OpenVMS Alpha Version 7.3 or higher. These devices include:

- Serial ports COM1 and COM2
- Parallel port
- Keyboard
- Mouse

If multiple legacy bus adapters exist, only the adapter that includes the console port is configured and supported.

6.6.2 OpenVMS Galaxy License Enforcement

V7.3

In an OpenVMS Galaxy computing environment, the OPENVMS-GALAXY license units are checked during system startup and whenever a CPU reassignment between instances occurs.

If you attempt to start a CPU and there are insufficient OPENVMS-GALAXY license units to support it, the CPU will remain in the instance's configured set but it will be stopped. You can subsequently load the appropriate license units and start the stopped CPU while the system is running. This is true of one or more CPUs.

Hardware Release Notes

6.6 AlphaServer GS Series Systems

6.6.3 License with Hard and Soft Partitions—Problem Corrected

V7.3-1

Some OpenVMS Alpha Version 7.2-1H1/Version 7.3 GS80/160/320 AlphaServer systems with hard or soft partitions and a common license database (LDB) have difficulty booting up all partitions if you had modified nonshareable license PAKs to share license units among OpenVMS instances.

OpenVMS Versions 7.3-2, 7.3-1, and 7.2-2 contain the correction to the problem; if you are running one of these versions, you can go directly to Section 6.6.5.

The following patch kits contain the fix to the problem:

- The patch kit for OpenVMS Version 7.2-1H1 can be found at:

http://ftp.support.compaq.com/patches/public/Readmes/vms/dec-axpvms-vms721h1_lmf-v0100--4.README

- The patch kit for OpenVMS Version 7.3 can be found at:

http://ftp.support.compaq.com/patches/public/Readmes/vms/dec-axpvms-vms73_lmf-v0100--4.README

Before you install any of these kits, you must perform the following steps:

1. If you installed workarounds, remove them. (See Section 6.6.4.)
2. Install licenses correctly. (See Section 6.6.5.)

6.6.4 Removing Workarounds

V7.3-1

OpenVMS systems that have a problem using a common license database on hard and soft partitions display the following symptoms:

- Booting OpenVMS produces OPCOM messages that state that licenses cannot be loaded.
- Logging in to OpenVMS produces an error message that states that the OPENVMS-ALPHA license cannot be loaded.

To work around the restriction on using nonshareable license PAKs in a common license database to share license units among OpenVMS instances on hard and soft partitions, HP recommended that you use one of the following two workarounds:

- Use a separate license database for each operating system instance running OpenVMS on AlphaServer GS80/160/320 systems.
- Use a temporary (120-day) license PAK provided by HP Services.

If you used either of these workarounds, you must remove it before installing the upgrade kit.

If You Set Up Separate License Databases:

If you had set up separate license databases, you should create a common license database before installing this kit. Refer to the *OpenVMS License Management Utility Manual* for guidelines on setting up a common license database.

If You Installed Temporary PAKs:

If you installed temporary PAKs in the common license database, you must disable the temporary PAKs for the following licenses:

OPENVMS-ALPHA
OPENVMS-ALPHA-USER

OPENVMS-ALPHA-ADL

To see a list of temporary PAKs installed on your system, enter the following command:

```
$ LICENSE LIST /FULL /BEFORE /TERMINATION_DATE="-120-" -  
_ $ OPENVMS-ALPHA, OPENVMS-ALPHA-USER, OPENVMS-ALPHA-ADL
```

This command displays (for the three licenses specified) the licenses that satisfy one of the following conditions:

- Have already terminated
- Will terminate within the next 120 days

Disable each of the temporary PAKs by using the LICENSE DISABLE command. For example:

```
$ LICENSE DISABLE OPENVMS-ALPHA /AUTHORIZATION=authorization-value
```

where *authorization-value* helps to identify the license, and is displayed by the LICENSE LIST command you used to find the temporary PAKs.

6.6.5 Installing Licenses

V7.3-1

Before you upgrade to Version 7.3-1 or higher, you should perform the following steps to ensure that the common license database can share license units among hard and soft partitions:

1. Calculate required units.
 - Load the base OpenVMS license.
 - Load the SMP licenses.
 - Use the following command to verify that you have the correct number of license units:

```
$ SHOW LICENSE /UNIT_REQUIREMENTS /CLUSTER
```

Note

The base OpenVMS license allows you to have only one interactive user login per physical system (not per partition). (However, you can always log in from OPA0: in each partition.) For additional interactive users, you will require additional license units. See your HP support representative to determine your needs.

2. Add your licenses to the common license database. For example:

```
$ LICENSE REGISTER license-name /ISSUER=DEC -  
_ $ /AUTHORIZATION=USA123456 -  
_ $ /PRODUCER=DEC -  
_ $ /UNITS=1050 -  
_ $ /AVAILABILITY=H -  
_ $ /OPTIONS=(NO_SHARE) -  
_ $ /CHECKSUM=2-BGON-IAMA-GNOL-AIKO
```

Note that you cannot use the /INCLUDE qualifier with the LICENSE REGISTER command to override the NO_SHARE attribute of the license.

Hardware Release Notes

6.6 AlphaServer GS Series Systems

3. Modify the license to override the NO_SHARE attribute of the PAKs with the command `LICENSE REGISTER /INCLUDE=(node-name-list)`. For example:

```
$ LICENSE MODIFY OPENVMS-ALPHA /INCLUDE=(NODEA, NODEB, NODEC)
```

4. To make OpenVMS Alpha license units available to the instance of OpenVMS running in each partition, you must ensure that SRM environment variable `SYS_SERIAL_NUM` is the same in each partition. To do so, perform the following steps:

- a. From the master console of each partition (usually on console line 0), use the `SHOW SYS_SERIAL_NUM` command to display the system serial number. For example:

```
P00>>>
P00>>>SHOW SYS_SERIAL_NUM
sys_serial_num G2A105
```

If the value of `SYS_SERIAL_NUM` is blank, use the `SHOW SYS_SERIAL_NUM` command from the master console in each of the other partitions to check for a nonblank system serial number.

Note

If all partition consoles show a blank value for `SYS_SERIAL_NUM`, you must create a nonzero value of up to 12 characters. Ensure that the system serial number that you create is not used on any other AlphaServer GS80/160/320 on this OpenVMS Cluster.

- b. Once you have determined the system serial number, use the `SET SYS_SERIAL_NUM` command from the master console of each partition to change `SYS_SERIAL_NUM` to the correct value. For example:

```
P00>>>
P00>>>SET SYS_SERIAL_NUM G2A105
```

You must do this in every hard partition and in every soft partition.

5. In order for the OpenVMS Cluster license database to be updated correctly, HP recommends that you completely shut down and reboot all OpenVMS Cluster common nodes. A rolling upgrade type of boot does not correctly update the common license database.

Note

If your system is part of an OpenVMS Cluster that shares a common license database, anytime you reconfigure the number of hard or soft partitions on your AlphaServer GS80/160/320, you must make sure that all partitions have the same `SYS_SERIAL_NUM`.

For partitionable machines that are sharing NO_SHARE licenses across partitions, it is possible to see the following error text on system bootup.

```
%LICENSE-E-NOAUTH, DEC OPENVMS-ALPHA use is not authorized on this node
-LICENSE-F-EXCEEDED, attempted usage exceeds active license limits
-LICENSE-I-SYSMGR, please see your system manager
Startup processing continuing...
```

This error text can be safely ignored. The text is displayed when someone has logged into a system that is sharing the OPENVMS-ALPHA PAK and they are then in use. This will be fixed in a future release.

6.6.6 AlphaServer GS140 Systems: Booting Problem Fixed

V7.3-2

If you boot an AlphaServer GS140 system with a Fibre Channel HSG80 system disk using Alpha Firmware Version 5.7, the system fails to reboot when the system is set to boot automatically. Enter the following command to show whether the system has been set to boot automatically:

```
SHOW AUTO_ACTION
```

This problem has been fixed with Alpha Firmware Version 6.4.

If you do not have Alpha Firmware Version 6.4 or later and you experience a failed power-cycle or INIT boot startup, enter a BOOT command at the console and press Return. OpenVMS shutdown and reboot commands will then perform as expected.

6.6.7 AlphaServer GS60/GS60E/GS140 Multiple I/O Port Module Configuration Restriction

V7.2-1

AlphaServer GS60/GS60E/GS140 configurations with more than a single I/O Port Module, KFTHA-AA or KFTIA-AA, might experience system failures.

When upgrading OpenVMS Galaxy and non-Galaxy AlphaServer 8200/8400 configurations with multiple I/O Port Modules to GS60/GS60E/GS140 systems, customers must install one minimum revision B02 KN7CG-AB EV6 CPU (E2063-DA/DB rev D01) module as described in Compaq Action Blitz # TD 2632.

For complete details about this restriction and its solution, refer to Compaq Action Blitz # TD 2632.

6.7 AlphaStation 200/400: ISA_CONFIG.DAT Changes Required

V7.1

Customers configuring ISA devices on AlphaStation 200/400 Family systems must change their SYS\$MANAGER:ISA_CONFIG.DAT file, so that the node information for each device appears at the end of each device description block.

Warning

For upgrades from OpenVMS Version 6.2 or 7.0 systems, this change must be made before starting the upgrade procedure.

Table 6–1 shows the changes to the device description block.

Table 6–1 Changes to Device Description Block

Before Version 7.1	After Version 7.1
[AUA0]	[AUA0]
NAME=AU	NAME=AU
NODE=3	DRIVE=SYS\$MSBDRIVER
DRIVER=SYS\$MSBDRIVER	IRQ=9
IRQ=9	DMA=(0,1)
DMA=(0,1)	PORT=(388:4,530:8)
PORT=(388:4,530:8)	NODE=3

Customers using SYS\$MANAGER:ISA_CONFIG.DAT files should read Section A.2.

6.8 AlphaStation 255: PCI Configuration Restriction

V7.1

The OpenVMS Alpha operating system does not support PCI option cards configured in PCI slot 0 on any AlphaStation 255 series systems.

PCI slot 0 is the lowest physical PCI option slot on AlphaStation 255 series systems. The interrupt signal for this slot is shared with the built-in Ethernet port. Because the OpenVMS Alpha operating system does not currently permit PCI devices to share an interrupt line, a PCI device installed in slot 0 will not function correctly or may cause errors to occur with the built-in Ethernet port. As a result of this restriction, AlphaStation 255 series systems support a maximum of two PCI option cards, configured in slot 1 and slot 2.

6.9 ATI RADEON 7500 Graphics

V7.3-2

The following notes describe resource requirements and a few restrictions for ATI RADEON 7500 graphics. Details about these new graphics features supported in OpenVMS Version 7.3-2 are included in the *HP OpenVMS Alpha Version 7.3-2 New Features and Documentation Overview* manual. You may also want to consult the HP DECwindows Motif for OpenVMS documentation set, in particular *Managing DECwindows Motif for OpenVMS Systems*. You can link to this document and others from the following URL:

<http://www.hp.com/go/openvms/doc>

6.9.1 Resource Requirements

Support for RADEON graphics requires the following system-wide resources:

- 2 global sections
- 296 KB of global memory

In addition, **each** RADEON card requires the following:

- 3 global sections
- 16 MB plus 1 page of global memory
- 16 MB plus 1 page of buffer object space (32-bit System Space Windows)

The global sections are charged against the GBLSECTIONS system resource, and the 16+ megabytes of global memory are charged against the GBLPAGES and GBLPAGFIL resources.

For example, a single RADEON card requires the following:

- 5 global sections
- 16 MB + 8 KB + 296 KB global memory

These numbers equate to the following values:

GBLSECTIONS	5
GBLPAGES	33376 (GBLPAGES is in units of 512-byte pagelets.)
GBLPAGFIL	2086 (GBLPAGFIL is in units of 8192-byte Alpha pages.)

A 4-card configuration requires the following:

- 14 global sections
- 296 KB + 4*16 MB + 4*8 KB = 64 MB + 328 KB global memory

These numbers equate to the following values:

GBLSECTIONS	14
GBLPAGES	131728 (GBLPAGES is in units of 512-byte pagelets.)
GBLPAGFIL	8233 (GBLPAGFIL is in units of 8192-byte Alpha pages.)

6.9.2 Support Limitations

The following functionality is not supported:

- S-Video output
- Digital output
- Dual-head operation

If you connect monitors to both the DVI port and the analog VGA (CRT) port, you will get identical video output on both screens. This is called **cloned video**. True dual-head operation, with independent video displays on each port, will be supported in a future release.

6.9.3 Message When AGP Card Not Found

When the DECwindows server initializes the cards, it looks for an AGP card first. If it does not find an AGP card, it outputs the following message in the server log file (SYS\$MANAGER:DECW\$SERVER_0_ERROR.LOG):

```
vmsInitDevice: Wrong device driver loaded, expected DVI WS type 57, found 56
RADEON RV200: AGP device not found. Trying PCI ...
```

This is a normal, informational message. It does not signal an error.

6.9.4 Cloned Video Restricted to Primary Head

Cloned video works only on the primary head.

Hardware Release Notes

6.9 ATI RADEON 7500 Graphics

6.9.5 Monitor Must Be Connected During Initialization

The RADEON 7500 graphics card has two video output ports: one for digital and one for analog. The digital interface is not currently supported. However, using a digital-to-analog adapter, you can plug an analog monitor into the digital port and get the identical output that you get using the analog port. If you use the digital port, the monitor must be attached before the system is powered up in order for the port to function correctly.

6.9.6 Boot Reset Recommendation

HP recommends that the console variable `boot_reset` be set to ON.

6.9.7 No Overlay Planes

Hardware overlay planes are not supported.

6.9.8 Single Colormap

The RADEON 7500 graphics controller supports only one hardware colormap. Keep this in mind if you change to the 8-bit color depth, where the default visual is PseudoColor. Attempting to use more than one PseudoColor colormap at a time causes colormap flashing.

Note

3D (OpenGL) rendering is not supported on 8-bit PseudoColor visuals.
(See also Section 6.9.14.)

Applications should not install or deinstall colormaps themselves. The window manager should perform these actions. However, the application is responsible for providing the window manager with hints about which colormaps to install or deinstall. You provide this information using the Xlib function `XsetWMColormapWindows()`. This function sets the `WM_COLORMAP_WINDOWS` property for a given window.

6.9.9 Single Bit Depth for All Windows

When you are using the RADEON 7500 card, all windows created on a particular head must have the same bit depth. The RADEON 7500 card supports bit depths of 8, 16, and 24 bits per pixel on any graphics head, but once the DECwindows server establishes a bit depth on a particular head, only windows or visuals with that bit depth can be created.

6.9.10 Pixel Depth for Read/Write Color Map

By default, the RADEON 7500 provides a pixel depth of 24 planes with a read-only, TrueColor color map. Some applications, such as DECwindows Paint, require a read/write color map. If Paint is run with a read-only color map, it fails with the following error message:

```
Error: Visual Not Supported
Supported Visuals are {PseudoColor, GrayScale, StaticGray}
```

To use a read/write color map, edit the file `SYSS$MANAGER:DECW$PRIVATE_SERVER_SETUP.COM` (or, if this file does not exist, create it from `SYSS$MANAGER:DECW$PRIVATE_SERVER_SETUP.TEMPLATE`) and add the following logical name definition to the file:

```
$ DEFINE/EXECUTIVE/SYSTEM/NOLOG DECW$SERVER_PIXEL_DEPTH 8,8,8,8,8,8,8,8
```

Then restart DECwindows using the following command:

```
$ @SYS$MANAGER:DECW$STARTUP RESTART
```

This change sets the pixel depth to 8 planes (on up to 8 graphics cards, which allows for a multiple-card configuration) and allows the server to provide a PseudoColor visual.

6.9.11 Backing Store/Save Unders Not Supported for 3D Windows

The implementation of backing store and save unders in the RADEON 7500 X server does not support 3D windows.

6.9.12 Threads Restriction

The RADEON 7500 OpenGL library for OpenVMS is not thread safe. However, OpenGL can be used in a multithreaded program if the use of OpenGL is restricted to a single thread within the program.

6.9.13 No Support for Single Buffered Visuals

The RADEON 7500 DECwindows server exports only double buffered visuals. For single buffering, users must select a double buffered visual and call `glDrawBuffer (GL_FRONT_BUFFER)` in their application.

6.9.14 No 3D Support for Color Index Mode

Even though 8-bit visuals are supported for 2D operations when the DECwindows server is started with bit depth = 8, OpenGL rendering is not supported on 8-bit visuals.

6.9.15 Timer Mechanism

Under certain circumstances, it is possible for a process to be interrupted while it owns the hardware lock, resulting in an apparent DECwindows server hang.

A timer mechanism has been implemented to detect this situation and to rectify it by forcing an image exit in the suspended process — or, in some instances, by deleting the process.

The timer mechanism can be configured using the following two logicals, which should be specified in the `DECW$PRIVATE_SERVER_SETUP.COM` file:

- `DECW$DRM_TIMER_PERIOD` (Default: 5.0 seconds)
Specifies the duration of the clock tick in seconds; a floating-point value.
- `DECW$DRM_TIMEOUT` (Default: 6)
Specifies the number of clock ticks that are allowed to elapse before a timeout occurs and the DECwindows server seizes control of the RADEON card.

The default values are chosen to minimize the impact of the timer on the performance of graphics applications. If you want to reduce the length of time before the DECwindows server begins responding again, you can do so by decreasing the value of `DECW$DRM_TIMER_PERIOD`.

A process can be interrupted while holding the hardware lock under either of the following conditions:

- The process is remotely logged in with its terminal displayed on a different system.

Hardware Release Notes

6.9 ATI RADEON 7500 Graphics

- The process is a subprocess that has been suspended or terminated by another process in such a way that normal exit handling does not occur.

If neither of these conditions is likely to occur in your configuration, you can disable the timer mechanism entirely by setting the timer period to zero:

```
$ DEFINE/SYSTEM DECW$DRM_TIMER_PERIOD 0
```

Whenever you change the value of DECW\$DRM_TIMER_PERIOD, you must either restart the DECwindows server or reboot the system for the changes to take effect. To restart the DECwindows server, use the following command:

```
$ @SYS$STARTUP:DECW$STARTUP RESTART
```

6.10 DECwindows X11 Display Server

This section contains release notes pertaining to the DECwindows X11 display server for OpenVMS Alpha systems.

6.10.1 S3 Multihead Graphics

Permanent Restriction

Alpha computers equipped with S3 Trio32 or Trio64 graphics cards support single-screen display only. Multihead graphics are not supported.

6.10.2 Integrated Graphics Board Support

V7.3-2

Starting with OpenVMS Version 7.3, support for graphics boards was fully integrated with the OpenVMS operating system. The following new boards are now supported:

- ATI RADEON 7500 (added in 2003)
- Elsa GLoria (PowerStorm 4D10T)
- OXYGEN VX1
- PowerStorm 300
- PowerStorm 350

6.11 DIGITAL Modular Computing Components (DMCC)

This section contains release notes pertaining to DMCC.

6.11.1 Alpha 5/366 and 5/433 PICMG SBC Restriction

Permanent Restriction

The KZPDA SCSI Controller and the PBXGA Graphics Card cannot be placed in a slot behind the bridge on the DIGITAL Modular Computing Components (DMCC) Alpha 5/366 and 5/433 PICMG SBCs.

6.11.2 Updating the SRM Console

Permanent Restriction

To update the SRM console on the Alpha 4/233 (21064a), 4/266 (21164a), 5/366, and 5/433 DMCC systems, you must choose either the SRM console or the AlphaBIOS setup. You can store only one console.

- If you are running OpenVMS on these systems, update only the SRM console.
- If you are running Windows NT on these systems, update only the AlphaBIOS setup.

If you accidentally update both the SRM and the AlphaBIOS consoles, you will enter the AlphaBIOS Setup menu, and you will *not* have the option to return to the SRM console. The only way to exit the AlphaBIOS Setup menu and return to the SRM console is to use a Firmware Update Utility located at the following Internet site:

<ftp://ftp.digital.com/pub/Digital/Alpha/firmware/index.html>

6.12 Digital Personal Workstation: Booting OpenVMS V7.3-1 and Higher

V7.3-1

If you are using the Digital Personal Workstation 433au, 500au, and 600au series systems, you can boot OpenVMS Version 7.3-1 or higher from an IDE CD-ROM if the controller chip is a Cypress PCI Peripheral Controller. You cannot boot OpenVMS on a Digital Personal Workstation au series system from an IDE CD-ROM with an Intel Saturn I/O (SIO) 82378 chip in your configuration. You must use a SCSI CD-ROM, if the Intel SIO chip is present.

To determine which IDE chip you have in your configuration, enter the following SRM console command:

SHOW CONFIGURATION

If you see *Cypress PCI Peripheral Controller*, you can boot OpenVMS.

If you see *Intel SIO 82378*, you will need to use and boot from a SCSI CD-ROM.

6.13 Dual-Controller HSGnn with Many LUNs Can Fail Under Heavy I/O Load

V7.3-2

A combination of improvements in driver performance and faster systems has uncovered a limit to the amount of I/O that a dual-controller HSGnn storage array configured with a relatively large number of LUNs can handle. When this limit is reached, it is possible for the array to be kept so busy processing I/O that it is unable to complete the normal periodic synchronization between controllers, causing a controller hang or failure and a loss of host access to some or all LUNs until a manual controller reset is performed. In the case of such a controller failure, the Last Failure Codes most likely to be reported on the HSG console are 01960186, 01942088, and 018600A0.

Hardware Release Notes

6.13 Dual-Controller HSGnn with Many LUNs Can Fail Under Heavy I/O Load

Most HSGnn devices will continue to run with no problems. If your site experiences an HSG controller hang or failure when a heavy load is applied and the HSG has more than approximately 24 LUNs configured, you may be able to avoid future hangs or failures by reconfiguring the controller with fewer LUNs or distributing I/O so that the HSG is not so heavily loaded.

This issue is being investigated by the appropriate HP engineering groups.

6.14 PowerStorm 300/350 PCI Graphics Support for OpenVMS

V7.3-2

For release notes on the PowerStorm 300/350 PCI graphics controller support for a Compaq Workstation running OpenVMS Alpha, refer to the *PowerStorm 300/350 OpenVMS Graphics Release Notes Version 2.0*. You can find these release notes on the OpenVMS Documentation CD-ROMs in the following directories:

Directory	File Name
[73.DOCUMENTATION.PS_TXT]	P300_350_REL_NOTES.PS,TEXT
[732.DOCUMENTATION.PS_TXT]	P300_350_V2_README.PS,TEXT

These documents, release notes, and installation guides are shipped with the graphics cards.

Obsolete Parameter Settings

Starting with OpenVMS Version 7.3, the following parameter settings are not relevant for PowerStorm 300 and 350 graphics cards: MAXBOBMEM, MAXBOBS0S1, and MAXBOBS2.

6.15 RFnn DSSI Disk Devices and Controller Memory Errors

V6.2

A problem exists with the microcode for earlier versions of RF31T, RF31T+, RF35, RF35+, RF73, and RF74 DSSI disk devices. The problem can cause data loss, and occurs when reading data from one of these devices, if the device has had a controller memory error (also known as an error detection and correction (EDC) error). The error could have been induced by a virtual circuit closure or faulty hardware.

HP advises customers with any of these devices to check their microcode revision levels. If the microcode revision levels are lower than the numbers shown in Table 6-2, HP recommends that you update the microcode.

The microcode for all models, except RF31T, RF31T+, and RF35+, is provided on the latest OpenVMS binary distribution CD-ROM.

The RF_VERS utility, a utility program that displays the microcode revision level of the DSSI disk devices, is also provided on the CD-ROM. Instructions both for using the utility program and for updating the microcode are provided in this section.

Hardware Release Notes

6.15 RFnn DSSI Disk Devices and Controller Memory Errors

Note

If you have an RF31T, RF31T+, or RF35+ disk drive with a version of microcode that is not supported (see Table 6–2), and if you have a support contract, contact your HP support representative. Otherwise, contact your authorized reseller.

The earliest supportable revision levels of the DSSI disk microcode are shown in Table 6–2.

Table 6–2 Supported Microcode Revision Levels

Device Type	Minimum Level with Supported Microcode
RF31T	T387E
RF31T+	T387E
RF35	T392D
RF35+	T392D
RF36	V427P
RF73	T392D
RF74	V427P

To display the microcode revision level of your DSSI disk devices, perform the following steps:

1. Log in to the SYSTEM account or another account that has the CMKRNL, DIAGNOSE, and SYSPRV privileges.
2. Enter the following commands:

```
$ SET PROCESS /PRIVILEGE=(DIAGNOSE,CMKRNL,SYSPRV)
$ SHOW DEVICE FYA0:
```

On VAX systems, if the SHOW DEVICE command produces an error, enter the following commands:

```
$ RUN SYS$SYSTEM:SYSGEN
SYSGEN> CONN FYA0/NOADAP
SYSGEN> ^Z
```

On Alpha systems, if the SHOW DEVICE command produces an error, enter the following commands:

```
$ RUN SYS$SYSTEM:SYSMAN
SYSMAN> IO CONNECT FYA0: /NOADAP
SYSMAN> ^Z
```

The following is an example of the display produced by the RF_VERS utility:

```
Program Name:  RF_VERS
Revision Level: V1.2s

NOTICE: This program does not currently support the RF72 or any
        HSDxx controllers. See next version for support.

DSSI disks currently on this system as seen by RF_VERS

Device      Node      Status      Hardware      Firmware
Name        Name                               Type          Version
```

Hardware Release Notes

6.15 RFnn DSSI Disk Devices and Controller Memory Errors

_22\$DIA7:	R4JL2I	mounted	RF73	T387A
_22\$DIA6:	R4I0BG	mounted	RF73	T387A
_22\$DIA8:	R4XLWE	mounted	RF73	T387A
_22\$DIA2:	R4FCZK	mounted	RF73	T387A
_22\$DIA3:	R4CKCG	mounted	RF73	T387A
_22\$DIA4:	R4ZKUE	mounted	RF73	T387A
_22\$DIA9:	R4GYI	mounted	RF73	T387A
_22\$DIA1:	R4XRYI	mounted	RF73	T387A

To update the microcode in your device, use the appropriate command for your device and platform from Table 6–3.

Caution

Back up the disk before updating the microcode.

Table 6–3 Commands for Updating Microcode in Certain DSSI Disk Devices

Device Type	Platform	Command
RF35	Alpha	\$RUN SYS\$ETC:RF35_T392F_DEC_ALPHA.EXE
RF35	VAX	\$RUN SYS\$ETC:RF35_T392F_DEC.EXE
RF36	Alpha	\$RUN SYS\$ETC:RF36_V427P_DEC_ALPHA.EXE
RF36	VAX	\$RUN SYS\$ETC:RF36_V427P_DEC.EXE
RF73	Alpha	\$RUN SYS\$ETC:RF73_T392F_DEC_ALPHA.EXE
RF73	VAX	\$RUN SYS\$ETC:RF73_T392F_DEC.EXE
RF74	Alpha	\$RUN SYS\$ETC:RF74_V427P_DEC_ALPHA.EXE
RF74	VAX	\$RUN SYS\$ETC:RF74_V427P_DEC.EXE

Caution

Do not delete SCSI_INFO.EXE, RF_VERS.EXE, or any of the files listed in Table 6–3. If these files are deleted, VMSKITBLD.COM (on VAX) will not be able to find them. Similarly, on Alpha systems, the PRODUCT INSTALL commands in AXPVMS\$PCSI_INSTALL and AXPVMS\$PCSI_INSTALL_MIN will fail.

6.16 RZnn Disk Drive Considerations

The following notes describe issues related to various RZ disk drives.

6.16.1 RZ25M and RZ26N Disk Drives: Recommendations

V7.1

During the testing of HP supported SCSI disk drives on configurations with DWZZAs and long differential SCSI buses, two drives, RZ25M and RZ26N, were found to have bus phase problems. For this reason, do not use these drives in configurations where the differential bus length connecting DWZZAs equals or exceeds 20 meters.

This recommendation applies only to the RZ25M and RZ26N drives. All other disk drives that are listed as supported in the OpenVMS SPD can be used in configurations to the full bus lengths of the SCSI-2 specification.

6.16.2 RZ26N and RZ28M Disks: Recommended Firmware Support

V6.2-1H3

The minimum firmware revision level recommended for RZ26N and RZ28M disks is Revision 0568.

If the latest firmware revision level is not used with these disks, multiple problems can occur.

6.16.3 RZ26L and RZ28 Disks: Required Firmware for Multihost Use

V6.2

If you install RZ26L or RZ28 disks on a multihost SCSI bus in an OpenVMS Cluster, the disk's minimum firmware revision is 442.

The following sections describe a procedure that you can use to update the firmware on some RZ26L and RZ28 drives. This procedure can only be used with drives that are directly connected to a SCSI adapter on a host system. Drives that are attached through an intelligent controller (such as an HSZ40 or KZPSC) cannot be updated using this procedure. Refer to the intelligent controller's documentation to determine whether an alternative firmware update procedure exists.

Important Note

Only certain RZ26L and RZ28 firmware revisions can be safely upgraded to firmware revision level 442. Refer to Section 6.16.3.1 to determine if your disks are capable of being upgraded to firmware revision level 442. If your disk is capable of supporting firmware revision level 442, use the RZTOOLS utility that is described in Section 6.16.3.2 to update the disk's firmware.

6.16.3.1 Firmware Revision Level 442 Requirements

Only the combinations of disk drives and firmware revision levels listed in Table 6–4 are capable of being upgraded safely to firmware revision level 442. Performing the update procedure on any other combination can permanently damage the disk.

Table 6–4 Revision Level 442 Firmware Compatibility

Disk Drive	Firmware Revision	Disk File Name
RZ26L	440C	RZ26L_442D_DEC.FUP
RZ28	441C or D41C 435 or 436	RZ28_442D_DEC2104.FUP RZ28P4_442C_DEC.FUP

6.16.3.2 Firmware Revision Level 442 Installation Procedure

If you determine that your disk requires revision level 442 firmware and it is capable of being upgraded safely, use the following procedure to update the firmware. (See Table 6–4 for the file name of the disk you are upgrading.)

Hardware Release Notes

6.16 RZnn Disk Drive Considerations

```
$ RZTOOLS_ALPHA := $SYS$ETC:RZTOOLS_ALPHA
$ RZTOOLS_ALPHA DKB500 /LOAD=$SYS$ETC:filename.FUP
  Read in 262144 bytes.
  Current FW version - X440C
  Upgrading to      - DEC0
  Loading code      .....
  New code has been sent to the drive.
```

6.17 ZLX Graphics Boards Support

Permanent Restriction

You must install Open3D for OpenVMS Alpha to support the following families of graphics boards:

- ZLX-M
- ZLX-L
- ZLXp-L

The latest version of Open3D for OpenVMS Alpha is V4.9B. To access the latest versions of Open3D for OpenVMS Alpha, check the Software Products Library at the following URL:

<http://www1.aclabs.com>

Click on “SPL master index” in the left sidebar. Then click on the date for OpenVMS Alpha under Current Software Products Library and search for Compaq Open3D in the list.

6.18 Recompiling and Relinking OpenVMS Device Drivers

The following sections contain release notes pertaining to recompiling and relinking OpenVMS device drivers.

6.18.1 Alpha and VAX SCSI Device Drivers

V7.3-1

All OpenVMS Alpha SCSI device drivers from previous versions of OpenVMS must be recompiled and relinked to run correctly on OpenVMS Version 7.3-1 or higher.

If you have an OpenVMS Alpha SCSI driver that you are upgrading from a version prior to OpenVMS Alpha 7.0, see Section 6.18.2.

Note that for OpenVMS Version 7.1 all OpenVMS VAX SCSI device drivers required recompiling and relinking. OpenVMS VAX device drivers that were recompiled and relinked to run on OpenVMS Version 7.1 will run correctly on OpenVMS Version 7.3 or higher.

6.18.2 OpenVMS Alpha Device Drivers

V7.1

Device drivers that were recompiled and relinked to run on OpenVMS Alpha Version 7.0 do not require source-code changes and do not have to be recompiled and relinked to run on OpenVMS Alpha Version 7.1 and higher. (Note that Alpha SCSI drivers, however, must be recompiled and relinked as described in Section 6.18.1.)

6.18 Recompiling and Relinking OpenVMS Device Drivers

Device drivers from releases prior to OpenVMS Alpha Version 7.0 that were not recompiled and relinked for OpenVMS Alpha Version 7.0 must be recompiled and relinked to run on OpenVMS Alpha Version 7.1 and higher.

OpenVMS Alpha Version 7.0 included significant changes to OpenVMS Alpha privileged interfaces and data structures. As a result of these changes, device drivers from releases prior to OpenVMS Alpha Version 7.0 may also require source-code changes to run correctly on OpenVMS Alpha Version 7.0 and higher. For more details about OpenVMS Alpha Version 7.0 changes that may require source changes to customer-written drivers, refer to the *OpenVMS Alpha Guide to Upgrading Privileged-Code Applications*.

6.19 Device Driver MON Version Handling

V7.3

As of OpenVMS Version 7.3, when `SYSTEM_CHECK` is enabled, device driver images with names of the form `SYS$nnDRIVER_MON.EXE` will be automatically loaded by the system loader. If a corresponding `_MON` version does not exist, the system will use the default image name: `SYS$nnDRIVER.EXE`.

6.20 Possible Per-Threads Security Impact on Alpha Device Drivers

V7.2

See Section 5.16.1 for information about how possible per-thread security impacts OpenVMS Alpha device drivers.

6.21 Device IPL Setup for OpenVMS Alpha Drivers

V6.2

Alpha hardware platforms that support PCI, EISA, and ISA buses deliver I/O device interrupts at different IPLs, either 20 or 21. The IPL at which device interrupts are delivered can change if you move the device from one platform to another. This is a problem if the driver declares its device IPL to be 20, and then that driver is executed on a machine that delivers I/O device interrupts at IPL 21.

The simplest solution to this problem is for PCI, EISA, and ISA device drivers to use IPL 21. This works correctly on platforms that deliver I/O device interrupts at IPL 20 and on platforms that deliver I/O device interrupts at IPL 21.

A future release of OpenVMS Alpha may provide a platform-independent mechanism for drivers to determine the device IPL dynamically.

6.22 CRCTX Routines Enhanced

V7.1-2

The system routines that you can use to manage the Counted Resource Context Block (CRCTX) data structure have been improved. The following routines now set and check the status (`CRCTX$V_ITEM_VALID`) of the CRCTX data structure:

- `IOC$DEALLOC_CRCTX`
- `IOC$ALLOC_CNT_RES`
- `IOC$DEALLOC_CNT_RES`
- `IOC$LOAD_MAP`

Hardware Release Notes

6.22 CRCTX Routines Enhanced

These routines have changed as follows:

If you call `IOC$DEALLOC_CRCTX` with a valid CRCTX status (`CRCTX$V_ITEM_VALID` set to 1), the service returns a bad status. If the `SYSBOOT` parameter `SYSTEM_CHECK` is set, the system will fail. This prevents users from deallocating a CRCTX when they have valid resources that have not been deallocated.

You must call `IOC$ALLOC_CNT_RES` with an invalid CRCTX status (`CRCTX$V_ITEM_VALID` set to 0). If you call this routine with a valid status, OpenVMS assumes that you will lose the resources mapped by this CRCTX. OpenVMS does not allocate new resources and returns a bad status. If `SYSTEM_CHECK` is set, the system will fail. `IOC$ALLOC_CNT_RES` sets the valid bit before it returns.

`IOC$DEALLOC_CNT_RES` must be called with a valid CRCTX (`CRCTX$V_ITEM_VALID` set to 1). If you call `IOC$DEALLOC_CNT_RES` with an invalid CRCTX, OpenVMS assumes that the other parameters are not valid, and returns a bad status. If `SYSTEM_CHECK` is set, the system will fail. `IOC$DEALLOC_CNT_RES` clears the valid bit before it returns.

`IOC$LOAD_MAP` must be called with a valid CRCTX. If it is called with an invalid CRCTX (`CRCTX$V_ITEM_VALID` set to 0), it assumes that the other parameters are also invalid, and returns a bad status. If the `SYSBOOT` parameter `SYSTEM_CHECK` is set, the system will fail.

These improvements indicate to device support and privileged-code application developers whether they need to deallocate scatter gather registers, which are treated by OpenVMS as generic resources. If the `CRCTX$V_ITEM_VALID` bit is set, `IOC$DEALLOC_CNT_RES` still needs to be called.

Product Retirement Notices

This appendix contains notifications about OpenVMS products that are no longer supported or that are slated for retirement. It also tells you how to find manuals that have been archived.

Freeware

Once a product is retired, HP does not accept or act on problem reports posted against the product. However, for those interested in doing their own development and support, the source code for many former products is available as freeware from the following sources:

- On the freeware CD-ROM that ships with the OpenVMS operating system. The freeware CD-ROM also includes internal tools such as SDL, NMAIL, MAILWATCH, and popular Internet programs. For instructions about mounting the CD-ROM, see Section 3.1.
- On the World Wide Web at the following address:

<http://h71000.www7.hp.com/openvms/freeware/>

A.1 Attunity Connect "On Platform" Package

V7.3-2

The Attunity Connect "On Platform" Package, part of the OpenVMS e-Business Infrastructure CD-ROM, will not be supported after May 31, 2004.

Attunity Connect continues to be available directly from Attunity for OpenVMS VAX and Alpha. For more information, see the Attunity web site:

<http://www.attunity.com/>

A.2 ISA_CONFIG.DAT Unsupported in Future Release

V7.1

Support for using the SYS\$MANAGER:ISA_CONFIG.DAT file to configure ISA devices will be discontinued in a future release of OpenVMS Alpha. If you use this file, you should convert to using the ISACFG utility from the console, and the new file-based autoconfiguration method for loading device drivers (as described in *Writing OpenVMS Alpha Device Drivers in C*).

Product Retirement Notices

A.3 POSIX 1003.4a Draft 4 Interface May Be Retired

A.3 POSIX 1003.4a Draft 4 Interface May Be Retired

V7.0

The POSIX 1003.4a, Draft 4 (or "d4") interface of the Compaq POSIX Threads Library (formerly named DECthreads) is slated for retirement in a future release. Applications that were written using the POSIX 1003.4a, Draft 4 interface should be migrated to the new POSIX 1003.1c standard (or "pthread") interface provided by the POSIX Threads Library. A compatibility mode for the Draft 4 POSIX 1003.4a interface has been provided in this release to help ease migration. This compatibility mode will be removed in a future release.

A.4 Archived Manuals

V7.3-1

As products are retired and the operating system evolves, certain OpenVMS manuals are archived. Archived manuals are no longer maintained and are not part of the OpenVMS documentation set. However, they are available on the OpenVMS Documentation CD-ROM and the following web site:

<http://www.hp.com/go/openvms/doc>

A.4.1 Extended File Specifications: Documentation Reorganization

V7.3-1

Information previously contained in the *OpenVMS Guide to Extended File Specifications* has been moved into other manuals and this manual is now obsolete.

The following table provides pointers to the new locations of topics from the archived manual:

Section	New Book	Location
Chapter 1 Overview of Extended File Specifications for OpenVMS		
All	<i>HP OpenVMS System Manager's Manual, Volume 1: Essentials</i>	Chapter 9
Chapter 2 Managing Extended File Naming on OpenVMS Systems		
All except Section 2.2.1, <i>Using RMS Default EFS Features</i>	<i>HP OpenVMS System Manager's Manual, Volume 1: Essentials</i>	Chapters 9 and 10
Section 2.2.1, <i>Using RMS Default EFS Features</i>	<i>Guide to OpenVMS File Applications</i>	Chapter 6
Chapter 3 Extended File Naming Characteristics		
All except Section 3.5, <i>DCL Commands and Utilities</i>	<i>OpenVMS User's Manual</i>	Chapter 5
Section 3.5, <i>DCL Commands and Utilities</i>	<i>HP OpenVMS DCL Dictionary</i>	Alphabetically by command

Product Retirement Notices

A.4 Archived Manuals

Section	New Book	Location
Chapter 4 Extended File Naming Considerations for OpenVMS Application Developers		
All	<i>OpenVMS Programming Concepts Manual, Volume II</i>	Chapter 29
Appendix A Setting Users' Expectations of Extended File Specifications		
Section A.1, ACP Interface (ODS-5)	<i>HP OpenVMS I/O User's Reference Manual</i>	Chapter 1
Section A.4, Restrictions	<i>Guide to OpenVMS File Applications</i>	Chapter 5
All except Sections A.1, A.4	<i>HP OpenVMS System Manager's Manual, Volume 1: Essentials</i>	Chapter 10
Appendix B Technical Information		
Section B.1, System Services Changes	<i>HP OpenVMS System Services Reference Manual</i>	Alphabetically by routine
Section B.2, Record Management Services (RMS) Changes	<i>Guide to OpenVMS File Applications</i>	See below
Section B.2.1, Record Management Services (RMS) Changes	<i>Guide to OpenVMS File Applications</i>	Chapter 5
Section B.2.2, Record Management Services (RMS) Changes (except Section B.2.2.7, DID Abbreviation and Section B.2.2.8, FID Abbreviation)	<i>Guide to OpenVMS File Applications</i>	Chapter 5
Section B.2.2.7 (DID Abbreviation)	<i>Guide to OpenVMS File Applications</i>	Chapter 6
Section B.2.2.8 (FID Abbreviation)	<i>Guide to OpenVMS File Applications</i>	Chapter 6
Section B.2.3, RMS Data Structure Changes (NAM Block)	<i>OpenVMS Record Management Services Reference Manual</i>	Chapter 5
Section B.2.3, RMS Data Structure Changes (NAML Block)	<i>OpenVMS Record Management Services Reference Manual</i>	Chapter 6
Section B.3, Files-11 XQP Changes	<i>HP OpenVMS I/O User's Reference Manual</i>	Chapter 1
Section B.4, Programming Utility Changes	<i>OpenVMS Utility Routines Manual</i>	Chapter 10
Section B.5, Run-Time Library Changes	<i>OpenVMS RTL Library (LIB\$) Manual</i>	Alphabetically by routine
Appendix C Character Sets		
All	<i>OpenVMS User's Manual</i>	Appendix A

Interlocked Memory Instructions

The *Alpha Architecture Reference Manual, Third Edition* (AARM) describes strict rules for using interlocked memory instructions. The Alpha 21264 (EV6) processor and all future Alpha processors are more stringent than their predecessors in their requirement that these rules be followed. As a result, code that has worked in the past, despite noncompliance, could fail when executed on systems featuring the 21264 processor and its successors. **Occurrences of these noncompliant code sequences are believed to be rare.** Note that the 21264 processor is not supported on versions prior to OpenVMS Alpha Version 7.1-2.

Noncompliant code can result in a loss of synchronization between processors when interprocessor locks are used, or can result in an infinite loop when an interlocked sequence always fails. Such behavior has occurred in some code sequences in programs compiled on old versions of the BLISS compiler, some versions of the MACRO-32 compiler and the MACRO-64 assembler, and in some HP C and C++ programs.

The affected code sequences use LDx_L/STx_C instructions, either directly in assembly language sources or in code generated by a compiler. Applications most likely to use interlocked instructions are complex, multithreaded applications or device drivers using highly optimized, hand-crafted locking and synchronization techniques.

B.1 Required Code Checks

OpenVMS recommends that code that will run on the 21264 processor be checked for these sequences. Particular attention should be paid to any code that does interprocess locking, multithreading, or interprocessor communication.

The SRM_CHECK tool has been developed to analyze Alpha executables for noncompliant code sequences. The tool detects sequences that may fail, reports any errors, and displays the machine code of the failing sequence.

B.2 Using the Code Analysis Tool (SRM_CHECK)

The SRM_CHECK tool can be found in the following location on the OpenVMS Alpha Version 7.3-2 Operating System CD-ROM:

```
SYS$SYSTEM:SRM_CHECK.EXE
```

To run the SRM_CHECK tool, define it as a foreign command (or use the DCL\$PATH mechanism) and invoke it with the name of the image to check. If a problem is found, the machine code is displayed and some image information is printed. The following example illustrates how to use the tool to analyze an image called myimage.exe:

```
$ define DCL$PATH []
$ srm_check myimage.exe
```

Interlocked Memory Instructions

B.2 Using the Code Analysis Tool (SRM_CHECK)

The tool supports wildcard searches. Use the following command line to initiate a wildcard search:

```
$ srm_check [...]* -log
```

Use the `-log` qualifier to generate a list of images that have been checked. You can use the `-output` qualifier to write the output to a data file. For example, the following command directs output to a file named `CHECK.DAT`:

```
$ srm_check 'file' -output check.dat
```

You can use the output from the tool to find the module that generated the sequence by looking in the image's MAP file. The addresses shown correspond directly to the addresses that can be found in the MAP file.

The following example illustrates the output from using the analysis tool on an image named `SYSTEM_SYNCHRONIZATION.EXE`:

```
** Potential Alpha Architecture Violation(s) found in file...
** Found an unexpected ldq at 00003618
0000360C  AD970130      ldq_l      R12, 0x130(R23)
00003610  4596000A      and        R12, R22, R10
00003614  F5400006      bne        R10, 00003630
00003618  A54B0000      ldq        R10, (R11)
Image Name:    SYSTEM_SYNCHRONIZATION
Image Ident:   X-3
Link Time:    5-NOV-1998 22:55:58.10
Build Ident:   X6P7-SSB-0000
Header Size:   584
Image Section: 0, vbn: 3, va: 0x0, flags: RESIDENT EXE (0x880)
```

The MAP file for `system_synchronization.exe` contains the following:

```
EXEC$NONPAGED_CODE 00000000 0000B317 0000B318 ( 45848.) 2 ** 5
SMPROUT            00000000 000047BB 000047BC ( 18364.) 2 ** 5
SMPINITIAL         000047C0 000061E7 00001A28 ( 6696.) 2 ** 5
```

The address 360C is in the `SMPROUT` module, which contains the addresses from 0-47BB. By looking at the machine code output from the module, you can locate the code and use the listing line number to identify the corresponding source code. If `SMPROUT` had a nonzero base, you would need to subtract the base from the address (360C in this case) to find the relative address in the listing file.

Note that the tool reports *potential* violations in its output. Although `SRM_CHECK` can normally identify a code section in an image by the section's attributes, it is possible for OpenVMS images to contain data sections with those same attributes. As a result, `SRM_CHECK` may scan data as if it were code, and occasionally, a block of data may look like a noncompliant code sequence. This circumstance is rare and can be detected by examining the MAP and listing files.

B.3 Noncompliant Code Characteristics

The areas of noncompliance detected by the `SRM_CHECK` tool can be grouped into the following four categories. Most of these can be fixed by recompiling with new compilers. In rare cases, the source code may need to be modified. See Section B.5 for information about compiler versions.

- Some versions of OpenVMS compilers introduce noncompliant code sequences during an optimization called "loop rotation." This problem can be triggered only in C or C++ programs that use `LDx_L/STx_C` instructions in assembly language code that is embedded in the C/C++ source using the `ASM` function,

Interlocked Memory Instructions

B.3 Noncompliant Code Characteristics

or in assembly language written in MACRO-32 or MACRO-64. In some cases, a branch was introduced between the LDx_L and STx_C instructions. This can be addressed by recompiling.

- Some code compiled with very old BLISS, MACRO-32, DEC Pascal, or DEC COBOL compilers may contain noncompliant sequences. Early versions of these compilers contained a code scheduling bug where a load was incorrectly scheduled after a load_locked.

This can be addressed by recompiling.

- In rare cases, the MACRO-32 compiler may generate a noncompliant code sequence for a BBSSI or BBCCI instruction where there are too few free registers.

This can be addressed by recompiling.

- Errors may be generated by incorrectly coded MACRO-64 or MACRO-32 and incorrectly coded assembly language embedded in C or C++ source using the ASM function.

This requires source code changes. The new MACRO-32 compiler flags noncompliant code at compile time.

If the SRM_CHECK tool finds a violation in an image, you should recompile the image with the appropriate compiler (see Section B.5). After recompiling, you should analyze the image again. If violations remain after recompiling, examine the source code to determine why the code scheduling violation exists. Then make the appropriate changes to the source code.

B.4 Coding Requirements

The *Alpha Architecture Reference Manual* describes how an atomic update of data between processors must be formed. The Third Edition, in particular, has much more information on this topic. This edition details the conventions of the interlocked memory sequence.

Exceptions to the following two requirements are the source of all known noncompliant code:

- There cannot be a memory operation (load or store) between the LDx_L (load locked) and STx_C (store conditional) instructions in an interlocked sequence.
- There cannot be a branch taken between an LDx_L and an STx_C instruction. Rather, execution must "fall through" from the LDx_L to the STx_C without taking a branch.

Any branch whose target is between an LDx_L and matching STx_C creates a noncompliant sequence. For instance, any branch to "label" in the following example would result in noncompliant code, regardless of whether the branch instruction itself was within or outside of the sequence:

```
LDx_L Rx, n(Ry)
...
label: ...
STx_C Rx, n(Ry)
```

Therefore, the SRM_CHECK tool looks for the following:

- Any memory operation (LDx/STx) between an LDx_L and an STx_C
- Any branch that has a destination between an LDx_L and an STx_C

Interlocked Memory Instructions

B.4 Coding Requirements

- STx_C instructions that do not have a preceding LDx_L instruction
This typically indicates that a backward branch is taken from an LDx_L to the STx_C. Note that hardware device drivers that do device mailbox writes are an exception. These drivers use the STx_C to write the mailbox. This condition is found only on early Alpha systems and not on PCI-based systems.
- Excessive instructions between an LDx_L and an STx_C
The AARM recommends that no more than 40 instructions appear between an LDx_L and an STx_C. In theory, more than 40 instructions can cause hardware interrupts to keep the sequence from completing. However, there are no known occurrences of this.

To illustrate, the following are examples of code flagged by SRM_CHECK.

```
** Found an unexpected ldq at 0008291C
00082914 AC300000 ldq_l      R1, (R16)
00082918 2284FFEC lda       R20, 0xFFEC(R4)
0008291C A6A20038 ldq       R21, 0x38(R2)
```

In the above example, an LDQ instruction was found after an LDQ_L before the matching STQ_C. The LDQ must be moved out of the sequence, either by recompiling or by source code changes. (See Section B.3.)

```
** Backward branch from 000405B0 to a STx_C sequence at 0004059C
00040598 C3E00003 br        R31, 000405A8
0004059C 47F20400 bis        R31, R18, R0
000405A0 B8100000 stl_c      R0, (R16)
000405A4 F4000003 bne        R0, 000405B4
000405A8 A8300000 ldl_l       R1, (R16)
000405AC 40310DA0 cmp_l     R1, R17, R0
000405B0 F41FFFFA bne        R0, 0004059C
```

In the above example, a branch was discovered between the LDx_L and STx_C. In this case, there is no "fall through" path between the LDx_L and STx_C, which the architecture requires.

Note

This branch backward from the LDx_L to the STx_C is characteristic of the noncompliant code introduced by the "loop rotation" optimization.

The following MACRO-32 source code demonstrates code where there is a "fall through" path, but this case is still noncompliant because of the potential branch and a memory reference in the lock sequence:

```
getlck: evax_ldql  r0, lockdata(r8) ; Get the lock data
        movl      index, r2        ; and the current index.
        tstl      r0               ; If the lock is zero,
        beql      is_clear         ; skip ahead to store.
        movl      r3, r2           ; Else, set special index.
is_clear:
        incl      r0               ; Increment lock count
        evax_stqc r0, lockdata(r8) ; and store it.
        tstl      r0               ; Did store succeed?
        beql      getlck           ; Retry if not.
```

To correct this code, the memory access to read the value of INDEX must first be moved outside the LDQ_L/STQ_C sequence. Next, the branch between the LDQ_L and STQ_C, to the label IS_CLEAR, must be eliminated. In this case, it could be done using a CMOVEQ instruction. The CMOVxx instructions are frequently useful for eliminating branches around simple value moves. The following example shows the corrected code:

```

        movl      index, r2          ; Get the current index
getlck: evax_ldql  r0, lockdata(r8)  ; and then the lock data.
        evax_cmoveq r0, r3, r2      ; If zero, use special index.
        incl      r0                 ; Increment lock count
        evax_stqc  r0, lockdata(r8) ; and store it.
        tstl      r0                 ; Did write succeed?
        beql      getlck             ; Retry if not.

```

B.5 Compiler Versions

Table B–1 contains information about versions of compilers that might generate noncompliant code sequences and the recommended minimum versions to use when you recompile.

Table B–1 Versions of OpenVMS Compilers

Old Version	Recommended Minimum Version
BLISS V1.1	BLISS V1.3
DEC Ada V3.5	HP Ada V3.5A
DEC C V5.x	DEC C V6.0
DEC C++ V5.x	DEC C++ V6.0
DEC COBOL V2.4, V2.5	DEC COBOL V2.6
DEC Pascal V5.0-2	DEC Pascal V5.1-11
MACRO–32 V3.0	V3.1 for OpenVMS Version 7.1-2 V4.1 for OpenVMS Version 7.2
MACRO–64 V1.2	See below.

Current versions of the MACRO–64 assembler might still encounter the loop rotation issue. However, MACRO–64 does not perform code optimization by default, and this problem occurs only when optimization is enabled. If SRM_CHECK indicates a noncompliant sequence in the MACRO–64 code, it should first be recompiled without optimization. If the sequence is still flagged when retested, the source code itself contains a noncompliant sequence that must be corrected.

Alpha computers with 21264 processors require strict adherence to the restrictions for interlocked memory sequences for the LDx_L and STx_C instructions described in the *Alpha Architecture Reference Manual, Third Edition*. To help ensure that uses of interlocked memory instructions conform to the architectural guidelines, additional checking has been added to Version 3.1 of the MACRO–32 Compiler for OpenVMS Alpha.

The *Alpha Architecture Reference Manual, Third Edition* describes the rules for instruction use within interlocked memory sequences in Section 4.2.4. The MACRO–32 for OpenVMS Alpha Version 3.1 compiler observes these rules in the code it generates from MACRO–32 source code. However, the compiler provides EVAX_LQxL and EVAX_STxC built-ins, which allow these instructions to be written directly in source code.

Interlocked Memory Instructions

B.5 Compiler Versions

The MACRO-32 Compiler for OpenVMS Alpha Version 4.1 now performs additional code checking and displays warning messages for noncompliant code sequences.

B.6 Recompiling Code with ALONONPAGED_INLINE or LAL_REMOVE_FIRST

Any MACRO-32 code on OpenVMS Alpha that invokes either the ALONONPAGED_INLINE or the LAL_REMOVE_FIRST macro from the SYS\$LIBRARY:LIB.MLB macro library must be recompiled on OpenVMS Version 7.2 or higher to obtain a correct version of these macros. The change to these macros corrects a potential synchronization problem that is more likely to be encountered on newer processors, starting with Alpha 21264 (EV6).

Note

Source modules that call the EXE\$ALONONPAGED routine (or any of its variants) do *not* need to be recompiled. These modules transparently use the correct version of the routine that is included in this release.

Index

A

ACME_SERVER
 multiple channels to RIGHTSLIST.DAT, 4-1

ACMS
 changes in error reporting, 2-1
 IPC-E-BCKTRNSFAIL error message, 4-2
 kits and file deletions, 1-7

Advanced Server for OpenVMS
 release note and upgrade requirements, 1-9
 upgrade path for PATHWORKS users, 1-10
 versions supported, 1-9

After-image journaling, 4-16

ALPHAbook 1, 6-1

AlphaServer 1000A
 BUS_PROBE_ALGORITHM default, 6-2
 installation failure with DEFPA adapter, 6-2

AlphaServer 2100
 console display, 6-2
 SCSI controller restriction, 6-3

AlphaServer 4100
 FRU table restriction, 6-3

AlphaServer 8200 systems
 FRU table restriction, 6-3

AlphaServer 8400 systems
 FRU table restriction, 6-3

AlphaServer ES47/ES80/GS1280 systems
 Firmware Version 6.6 recommended, 6-3, 6-4
 license requirement, 6-4
 RAD support, 6-4
 setting time at MBM, 6-5
 STOP/CPU and shutdown behavior, 6-4
 TKNOVF error on F\$GETSYI item codes, 6-4

AlphaServer GS Series systems, 4-20
 device restriction, 6-5
 multiple I/O port restriction, 6-9

AlphaServer GS1280 systems
 ERLBUFFERPAGES, 6-5

AlphaServer GS140 systems
 booting, 6-9

AlphaStation 200/400
 ISA_CONFIG.DAT changes required, 6-9

AlphaStation 255
 PCI configuration restriction, 6-10

ANALYZE/DISK/SHADOW command, 4-26
 and dissimilar device shadowing, 4-27

Applications support for current release, 2-1

Archived manuals, A-2
 OpenVMS Guide to Extended File Specifications,
 A-2

Associated products
 Software Public Rollout Reports, 2-1
 versions supported for current release, 2-1

ATI RADEON 7500 graphics, 6-10 to 6-14

Attunity Connect "On Platform" Package, A-1

AUTOGEN
 and NEWPARAMS.DAT files, 4-1

B

Backup API
 journaling events, 5-1

BASIC
 V1.5A required to create STARLET library,
 2-2

BCKTRNSFAIL error message, 4-2

BLISS compiler, consequences of noncompliant
code, B-1

BUGCHECKFATAL system parameter, 5-14

BUS_PROBE_ALGORITHM settings, 6-2

C

C programs
 compiling and case sensitivity, 5-1

C RTL, 5-1 to 5-12
 64-bit pointers supported for some socket
 routines, 5-2
 angle brackets in file specs, 5-4
 BSD missing entry points, 5-8
 catopen supports UNIX style path names, 5-3
 compiler errors, 5-8
 DECC\$ALLOW_REMOVE_OPEN_FILES
 feature logical controls removal of open
 files, 5-5
 DECC\$RENAME_ALLOW_DIR feature logical,
 5-9
 dlsym problems fixed, 5-3
 exec error, 5-8
 FILE_TYPE macro removed from <ftw.h>
 header, 5-2
 fstat function checks for enabled disk access
 dates, 5-5

C RTL (cont'd)

- <ftw.h> header compatible with ftw function, 5-2
 - fwrite problem fixed, 5-2
 - GETADDRINFO undefined symbol fixed, 5-1
 - glob note, 5-10
 - iconv problem, 5-3
 - inet_ntop, 5-8
 - IPv6 programming note, 5-10
 - IPv6 symbol names, 5-9
 - memory management functions now reentrant, 5-4
 - mixed logical name search lists, 5-7
 - pread and pwrite behavior change for unseekable files, 5-5
 - program deadlock problem, 5-8
 - rename function, 5-9
 - sleep and usleep race condition fixed, 5-4
 - timespec structure redefinition, 5-5
 - timespec_t typedef redefinition, 5-5
 - tm_t structure removed from <time.h> header, 5-2
 - tzset behavior change for invalid TZ, 5-4
 - UNIX file name translation, 5-9
 - UNIX to OpenVMS name translation errors fixed, 5-4
- ## C Run-Time Library
- See C RTL
- ## CANCEL SELECTIVE function, improved use with LTDRIVER, 5-19
- ## CDSA, 5-12
- ## Circuit switching
- and reduced performance, 4-11
- ## CLISYMTBL system parameter, 2-7
- ## Cluster compatibility kits, 4-6
- ## Clusters
- See OpenVMS Cluster systems
- ## CLUSTER_CONFIG.COM
- limits on root directory names, 4-9
- ## COBOL
- START or WRITE statement, 5-14
- ## COM for OpenVMS, 2-3
- error with heavy load of applications, 2-3
- ## NTA\$LOGON
- batch login error, 2-3
 - /TYPE qualifier behavior changes, 2-3
 - upgrading, 1-7
- ## Common Data Security Architecture
- See CDSA
- ## Compilers
- noncompliant code, B-1, B-5
- ## CONVERT-I-SEQ error
- CONVERT/NOSORT, 5-13
- ## COPY command
- performance enhancement, 3-2

- CPUSPINWAIT bugcheck, 5-14
- CRCTX routines enhanced, 6-21

D

- ## DCL command procedures
- using SET SHADOW and SHOW SHADOW, 4-22
- ## Debugger
- client/server interface support, 5-13
- ## Debugging using POSIX Threads Library, 5-21
- ## DEC PL/I, 2-4
- ## DECdfs for OpenVMS
- Version 2.3-3 required, 2-3
- ## DECdtm
- IPC-E-BCKTRNSFAIL error message, 4-2
 - Oracle 8i, 9i, 4-2
- ## DECevent
- enabling the DIAGNOSE command, 1-7
- ## DECforms Web Connector
- running on OpenVMS Version 7.3-1 and higher, 2-4
- ## DECnet for OpenVMS, 1-4
- ## DECnet-Plus for OpenVMS, 1-4
- new version required, 1-8
- ## DECnet/OSI
- See DECnet-Plus for OpenVMS
- ## DECram
- See HP DECram
- ## DECwindows Motif
- language variant availability, 2-7
 - new version required, 1-8
 - required system parameter values, 2-7
 - support for LAT transport interface, 2-6
 - user-written transport support, 2-7
- ## DECwindows X11 display server, 6-14
- graphics boards support, 6-20
- ## DEFPA adapter
- on AlphaServer 1000A computer, 6-2
- ## DEGPA device
- error with UETP, 4-21
- ## Device drivers
- IPL setup, 6-21
 - MON version handling, 6-21
 - per-thread security impact, 6-21
 - recompiling and relinking, 6-20 to 6-21
 - SCSI, 6-20
- ## DEVOFFLINE error
- on ANALYZE/DISK/SHADOW command, 4-26
- ## DIAGNOSE command, enabling, 1-7
- ## DIGITAL Modular Computing Components (DMCC)
- KZPDA controller and PBXGA Graphics Card, 6-14
 - updating the SRM console, 6-15

- Digital Personal Workstation, 6–15
- DIRECTORY command
 - changes to output, 4–17
- Dissimilar device shadowing (DDS)
 - ANALYZE/DISK/SHADOW command behavior, 4–27
 - KZPDC restriction, 4–24
 - restriction on adding first dissimilar member, 4–23
 - Smart Array 5300 restriction, 4–24
 - write bitmap interaction, 4–23
- Documentation changes and corrections
 - archived manuals, A–2
 - HP OpenVMS Alpha Version 7.3–2 Upgrade and Installation Manual*, 1–4
 - HP OpenVMS System Manager's Manual*, 4–6
 - OpenVMS Cluster Systems manual*, 4–9
 - OpenVMS RTL Library (LIB\$) Manual*, 5–23
 - OpenVMS RTL Screen Management (SMG\$) Manual*, 5–24
- DSSI disk devices
 - microcode revision levels, 6–17
- Dual-controller HSGnn
 - failure, 6–15
- Dynamic CPU configuration
 - POSIX Threads Library, 5–21

E

- ECP (Enterprise Capacity and Performance)
 - minimum version, 4–3
- EDIT/FDL
 - fixing recommended bucket size, 4–3
- ELV
 - See Error Log Viewer (ELV) utility
- Enterprise Capacity and Performance (ECP)
 - See ECP
- Error Log Viewer (ELV) utility, 4–3
- EV6 Alpha processor, B–1
- Extended DDT bit, problem corrected, 5–19
- Extended File Cache (XFC), 4–4
- Extended File Specifications
 - documentation reorganized, A–2
- External authentication
 - password expiration notification, 4–5
 - requirements, 4–4
 - SET PASSWORD command, 4–4

F

- F\$GETSYI item codes
 - TKNOVF error, 6–4
- Fast lock remastering, 4–5
- Fast Path
 - turning off, for Galaxy on ES40, 4–14

- FDDI satellite booting problem, 4–10
- Fibre Channel
 - compatibility kits, 4–6
 - deleting a device, 4–26
 - multipath failover restriction for tape devices, 4–12
- Firmware, 1–3
- FMS ECO kit, 2–4
- Freeware, 3–1, A–1
- FREE_GBLPAGES system parameter, 2–7

G

- GBLPAGES system parameter, 2–7
- Gigabit Ethernet switch restriction, 4–12
- Graphical Configuration Manager (GCM)
 - not supported for Galaxy, 4–13
- Graphical Configuration Utility (GCU), 4–13
- Graphics boards support, 6–20

H

- Help
 - topic name changes, 3–2
- HP C compiler
 - consequences of noncompliant code, B–1
- HP C++ compiler
 - consequences of noncompliant code, B–1
- HP COBOL RTL, 5–14
- HP DECram, 2–5
 - maximum disk size, 2–5
 - Version 2.5, 2–5
 - Version 3.2 required, 2–5
- HP Secure Web Browser
 - increased memory requirement, 3–2
- HP SSL installation, 1–14
- HSGnn
 - failure, 6–15
- Hypersort, 5–15 to 5–16

I

- IDE CD-ROM, 6–15
- INITIALIZE command
 - incorrect messages output, 4–5
- Installation and upgrade information
 - enabling the DIAGNOSE command, 1–7
 - networking options, 1–4
 - PCSI-I-RETAIN messages, 1–9
- Integrated graphics boards, 6–14
- Interlocked memory instructions, B–1
- IPC-E-BCKTRNSFAIL error message, 4–2
- IPL Requirement
 - VMS fork thread creation, 5–23
- IPv6
 - programming note, 5–10
 - symbol names, 5–9

ISA_CONFIG.DAT
 unsupported in future release, A-1

K

Kerberos, 1-11
 removing V1.0 before upgrade, 1-11
Kernel threads
 incompatibility with recovery unit journaling,
 4-15
Kits
 compressed format, 1-6
 converting to uncompressed format, 1-6

L

LDAP API problem, 5-17
LIB\$FIND_IMAGE_SYMBOL
 LIB\$EOMWARN warning, 5-24
 passing by value, 5-23
Librarian utility (LIBRARIAN)
 error reporting problem and workaround, 5-16
Licensing issues, 6-5 to 6-9
Linker, 5-17 to 5-19
 change in behavior with library check, 5-19
 hangs when processing many files, 5-17
 limit of 25 elements on stack, 5-19
 RMS_RELATED_CONTEXT option, 5-17
Lock Manager, 4-5
Logical Disk (LD) utility
 error when using RMS, 4-6
LTDRIVER restriction, 5-19
LUNs
 autoconfiguration failure, 4-18

M

MACRO-32 compiler
 consequences of noncompliant code, B-1
 recompiling code, B-6
MACRO-64 assembler
 consequences of noncompliant code, B-1
Mail utility (MAIL)
 documentation correction, 4-6
 problem when callable mail used with kernel
 threads, 5-19
MEDOFL error
 on ANALYZE/DISK/SHADOW command, 4-26
 on tape after dismount, 4-9
Microcode revision levels
 commands for updating, 6-18
 on DSSI disk devices, 6-17
MultiNet, 2-8
Multipath failover
 Fibre Channel tape device restriction, 4-12
 tape robots, 4-13

MULTIPROCESSING system parameter, 5-14

N

Network options, 1-4
NEWPARAMS.DAT files
 and AUTOGEN, 4-1
NISCS_LAN_OVRHD system parameter, 4-10
NISCS_MAX_PKTSZ system parameter, 4-10
Noncompliant code, B-1, B-2
NPAGEDYN system parameter, 4-20
NPAGERAD system parameter, 4-20
NTA\$LOGON
 batch login error, 2-3
 /TYPE qualifier behavior changes, 2-3

O

Online help
 See Help
Open3D graphics boards support, 6-20
OpenVMS Cluster systems, 4-6
 compatibility kits, 4-6
 compatibility kits for mixed versions, 4-6
 Gigabit Ethernet switch restriction, 4-12
 limits on root directory names, 4-9
 MEDOFL error after tape dismount, 4-9
 performance reduced with CI-LAN switching,
 4-11
 rolling upgrades, 1-3
 correction to *HP OpenVMS Alpha Version
 7.3-2 Upgrade and Installation
 Manual*, 1-4
 SCSI multipath failover, 4-8
OpenVMS Galaxy, 4-13 to 4-14
 and ES40
 turning off Fast Path, 4-14
 uncompressed dump limitation, 4-14
 DECram support, 2-5
 error on call to SYS\$CRMPSC_GDZRO_64,
 4-14
 license enforcement, 6-5
 unexpected SS\$_INSFMEM status, 4-14
OpenVMS Management Station, 4-14
OpenVMS Registry
 Version 2 format database corruption, 4-15
OpenVMS upgrade
 systems with Advanced Server, 1-10

P

Pascal
 reinstalling after an upgrade (Alpha), 2-8
 V5.8A required to create STARLET library,
 2-8
Patch kits
 required for mixed-version OpenVMS Cluster
 system, 4-6

- PATHWORKS for OpenVMS (Advanced Server)
 - displaying ACEs, 4-17
 - release note and upgrade requirements, 1-10
 - upgrade path, 1-10
 - for early versions, 1-10
 - versions supported, 1-10
- PATHWORKS V5 for OpenVMS
 - not supported, 1-10
- PCI configuration restriction, 6-10
- PCSI-I-RETAIN message, 1-9
- PE1 system parameter, 4-5
- PEdriver
 - error message change, 4-10
 - response to LAN congestion, 4-11
 - with channel priority of -128, 4-11
- Per-thread security
 - impact on device drivers, 5-22
 - impact on privileged code, 5-22
- PGFLQUOTA problems, 5-16
- PL/I RTL, 2-4
- POOLCHECK system parameter, 5-14
- Port driver \$QIO
 - restriction, 5-19
- POSIX Threads Library, 5-19 to 5-22
 - debugger metering does not work, 5-22
 - Dynamic CPU configuration, 5-21
 - POSIX 1003.4a, Draft 4 interface to be retired, A-2
- PowerStorm 300/350 PCI graphics support, 6-16

R

- Rdb
 - IPC-E-BCKTRNSFAIL error message, 4-2
- Record locks and COBOL, 5-14
- Recovery unit journaling
 - file creation changes, 4-15
 - kernel threads, 4-15
 - remote access of journaled files, 4-16
 - restriction, 4-16
- Registry
 - considerations when upgrading, 1-12
- Remedial kits
 - obtaining, 1-5
 - required for mixed-version OpenVMS Cluster system, 4-6
- Retired products information, A-1
- RF73 and RFnn disks, controller memory errors, 6-16
- RMS Journaling, 4-15
 - after-image journaling, 4-16
 - journal file creation modification, 4-15
 - remote access of recovery unit journal files, 4-16
- RU journaling
 - See Recovery unit journaling

- Run-time library (LIB\$), 5-24
- RZnn disk drives, 6-18 to 6-20

S

- Satellite booting
 - problem over FDDI, 4-10
- SCSI controllers
 - restrictions on AlphaServer 2100 systems, 6-3
- SCSI device drivers, 6-20
- SCSI multipath incompatibility, 4-8
- SCSNODE system parameter
 - size enforced, 4-20
- Security auditing fixed, 4-17
- Server Management Process (SMHANDLER), 4-17
- SET DEVICE/SWITCH command, 4-13
- SET PASSWORD command, 4-4
- SET SHADOW
 - using in DCL command procedures, 4-22
- SHADOW_MAX_UNIT
 - default setting, 1-5
- SHAD_MAX_UNIT
 - memory consumption, 1-5
- SHOW DEVICE/BITMAP command
 - INTDIV error, 4-25
 - LOG_IO privilege required, 4-25
- SHOW SHADOW
 - using in DCL command procedures, 4-22
- Smart Array 5300
 - Volume Shadowing restriction, 4-24
- Smart Array 5300 restrictions, 4-13
- SMG\$
 - documentation corrections, 5-24
- SMHANDLER
 - See Server Management Process
- Software Public Rollout Reports, 2-1
- SORT32 utility, 5-25 to 5-26
- SPLINVIPL bugcheck, 1-11, 5-23
- SRM_CHECK tool, B-1
- SSL installation, 1-14
- START statement and COBOL, 5-14
- Support policy for software, 1-1
- SYSGEN
 - security auditing fixed, 4-17
- SYSMAN
 - DUMP_PRIORITY LOAD command, 4-18
 - IO AUTO command
 - XP LUNS fail to configure, 4-18
- System Event Analyzer (SEA) utility, 1-15
- System parameters, 4-19 to 4-20
 - BUGCHECKFATAL, 5-14
 - CLISYMTBL, 2-7
 - FREE_GBLPAGES, 2-7
 - GBLPAGES, 2-7
 - modified parameters, 4-19
 - MULTIPROCESSING, 5-14
 - new parameters, 4-19

System parameters (cont'd)

- NISCS_LAN_OVRHD, 4–10
- NISCS_MAX_PKTSZ, 4–10
- NPAGEDYN, 4–20
- NPAGERAD, 4–20
- obsolete parameters, 4–20
- PE1, 4–5
- PHYSICAL_MEMORY, 6–4
- POOLCHECK, 5–14
- SCSNODE size enforced, 4–20
- SYSTEM_CHECK, 5–14

System services

- SYS\$ACM(W) changes, 5–26

System startup failure message, 1–5

- SYSTEM_CHECK system parameter, 5–14

T

Tape robots

- automatic multipath failover, 4–13

TCP/IP Services for OpenVMS, 1–4

- Terminal Fallback Facility (TFF), 4–20
- restrictions, 4–21

TFF

- See Terminal Fallback Facility

Timer Queue Entries, 5–26

- TKNOVF error on F\$GETSYI item codes, 6–4

TQEs

- See Timer Queue Entries

U

UETP (User Environment Test Package)

- error with DEGPA device, 4–21

Upgrade paths, 1–3

- correction to *HP OpenVMS Alpha Version 7.3–2 Upgrade and Installation Manual*, 1–4

V

VFC format sequential files, 4–16

Visual Threads, debugging with, 5–21

VMSINSTAL

- VMSINSTAL_NOBROADCAST, 4–22

Volume Shadowing for OpenVMS, 4–22 to 4–28

- ANALYZE/DISK/SHADOW command, 4–26
- and DDS, 4–27

- compatibility kits, 4–6

- deleting a Fibre Channel device, 4–26

- device name requirement, 4–22

- dissimilar device shadowing (DDS)

- caution, 4–23

- KZPDC restriction, 4–24

- INTDIV error on SHOW DEVICE/BITMAP command, 4–25

- performance of merge operations, 4–26

- problem with dismount using /MINICOPY, 4–27

Volume Shadowing for OpenVMS (cont'd)

- Smart Array 5300 (KZPDC) restriction, 4–24

W

Write bitmaps

- and dissimilar device shadowing (DDS), 4–23

- WRITE statement and COBOL, 5–14

X

X.25

- upgrade, 1–11

XA, 4–2

XFC

- See Extended File Cache

XP

- LUNs fail to configure, 4–18

Z

- ZLX graphics boards support, 6–20