
HP OpenVMS Version 8.3–1H1 for Integrity Servers New Features and Release Notes

Order Number: BA322-90076

October 2007

This manual describes the new features and release notes for OpenVMS Version 8.3–1H1 for Integrity servers.

Revision/Update Information: This is a new manual.
Software Version: OpenVMS Version 8.3–1H1 for Integrity servers

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

© Copyright 2007 Hewlett-Packard Development Company, L.P.

Confidential 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.

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.

Intel and Itanium are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the United States and other countries.

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, and Windows NT 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.

ZK6676

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

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

Contents

Preface	ix
1 Introduction	
1.1 Overview of OpenVMS Version 8.3–1H1 for Integrity Servers	1–1
1.2 HP Software Technical Support Policy	1–6
1.3 General Application Compatibility Statement	1–6
1.4 Obtaining Patch Kits	1–7
1.5 Networking Options	1–7
1.6 System Event Log on Integrity Servers	1–7
1.7 Firmware for Integrity Servers	1–8
1.8 Release Notes on Booting the System	1–9
1.8.1 Booting from the Installation DVD	1–10
1.8.2 Setting Up I64 Systems to Reboot	1–10
1.8.3 Booting with a Common Cluster System Disk	1–10
1.8.4 Booting from a Fibre Channel Storage Device	1–10
1.8.5 OpenVMS I64 Boot Manager Utility: Adding Multipath Fibre Channel Disk Devices	1–10
1.8.6 Fibre Channel Boot Disk: Simplified Setup Process	1–11
1.9 HP DECwindows Motif	1–11
1.9.1 Connect Peripheral Devices Prior to Server Startup	1–11
1.9.2 Countdown Messages Displayed During Startup	1–12
1.9.3 Optional Graphics	1–12
1.9.4 Keyboard Support	1–12
1.9.5 Firmware Update and Keyboards	1–12
2 OpenVMS Version 8.3–1H1 New Features	
2.1 ISV Applications and Binary Compatibility in HP OpenVMS Version 8.3-1H1	2–1
2.2 New Item Code DVIS_ADAPTER_IDENT	2–1
2.3 New item code DVIS_MOUNTCNT_CLUSTER	2–1
2.4 HP Smart Array P800 Controller (AD335A) Support	2–2
2.5 HP 4GB Fibre Channel Adapters (AD299A/AD355A) Support	2–2
2.6 HP StorageWorks Ultrium 448c Tape Blade	2–2
2.7 Storage and Network I/O Controllers supported after the Initial OpenVMS V8.3 Release	2–2
2.8 Graphics Console Support for Selected HP Integrity Server Platforms ...	2–2
2.8.1 Graphics Console Support Platforms	2–3
2.8.2 Supported PCI Graphics Option Cards	2–3
2.8.3 Graphics Console Limitations	2–3

2.8.4	Minimum Firmware Required for Graphics Console Support	2-4
2.8.4.1	Obtaining Firmware Information from the EFI Shell	2-4
2.8.4.2	Obtaining Firmware Information from the MP Prompt	2-4
2.8.4.3	Obtaining Firmware Information from the cli Prompt	2-5
2.8.4.4	Obtaining Firmware Information from the EFI Boot Manager	2-5
2.8.5	Platform-Specific Connection Details	2-5
2.8.5.1	HP Integrity rx1600 and rx1620 Servers	2-6
2.8.5.2	HP Integrity rx2620 and rx4640 Servers	2-6
2.8.5.3	HP Integrity rx3600 and HP Integrity rx6600 Servers	2-6
2.8.5.4	HP Integrity rx2660 Servers	2-7
2.8.5.5	HP Integrity BL860c Server Blade	2-7
2.9	OpenVMS Internationalization Data Kit (VMSI18N)	2-7
2.10	LDAP External Authentication	2-8
2.10.1	Enhanced LOGINOUT.EXE and SETP0.EXE Images for LDAP External Authentication	2-8
2.10.2	Active Directory Support	2-8
2.11	Linker New Features	2-8
2.11.1	Support for Larger I/O Segments	2-8
2.11.2	Optimizing Linker Generated Code	2-8
2.11.3	Signature Segment Flagged in Map File	2-9
2.11.4	Unwind Segments Follow Code Segments	2-9
2.12	System Dump Analyzer (SDA) Utility	2-9
2.12.1	Changes to CLUE Utility	2-9
2.12.2	CLUE Utility: New Qualifier	2-9
2.12.3	SHOW_PFN_DATA Command: New Qualifier	2-9
2.12.4	SHOW SUMMARY Command: New Qualifier	2-9
2.12.5	SET AND SHOW CPU Commands	2-10
2.13	Time Zones Added	2-10

3 OpenVMS Version 8.3-1H1 Release Notes

3.1	OpenVMS Version 8.3-1H1 System Parameters	3-1
3.1.1	Obsolete Parameters	3-1
3.1.2	Renamed Parameters	3-1
3.1.3	New Parameter	3-2
3.1.4	Failure of AUDIT_SERVER to Initiate During Boot	3-2
3.2	C Run-Time Library	3-2
3.2.1	Ctrl/P Does Not Always Work	3-4
3.3	Encryption Issues and Resolutions	3-5
3.3.1	PARSE=EXTENDED	3-5
3.3.2	Large File Encryption	3-6
3.3.3	Input/Output File Performance	3-6
3.3.4	ENCYPT/COMPRESS Issue	3-6
3.4	Documentation Corrections	3-6
3.4.1	HP OpenVMS DCL Dictionary for OpenVMS: N-Z	3-6
3.4.1.1	RUN Command Restriction	3-7
3.4.2	HP OpenVMS Version 8.3 New Features and Documentation Overview	3-7
3.4.3	HP OpenVMS Version 8.3 Release Notes	3-7
3.4.4	HP OpenVMS System Services	3-8
3.4.5	HP TCP/IP Services for OpenVMS Sockets API and System Services Programming Manual	3-8
3.4.6	HP OpenVMS System Management Utilities Reference Manual	3-8
3.4.7	HP TCP/IP Services for OpenVMS Management	3-9

3.5	GiCAP Functionality Not Currently Available	3-9
3.6	Iconv Fixes	3-9
3.7	No-IOLOCK8 Fibre Channel Port Drivers	3-9
3.8	Linker Utility for OpenVMS I64	3-10
3.8.1	Linker Writes Incorrect Interimage Debug Fixups into Debug Symbol File	3-10
3.8.2	/SELECTIVE_SEARCH Qualifier Might Incorrectly Ignore Transfer Address	3-11
3.8.3	Maximum Number of Sections	3-12
3.8.4	Incorrect Creation Date of Shareable Images in the Map File	3-12
3.8.5	Demangler Information Look Up Results in Linker Access Violation	3-12
3.8.6	Incorrect Secondary Messages for the NOGLOSYM Error Message	3-12
3.8.7	Incorrect Information for Undefined Symbols	3-12
3.8.8	Incorrect UNMAPFIL Error	3-13
3.8.9	Max Ident Length Change for Shareable Images in Map	3-13
3.8.10	Linkage Type Check for Shareable Images	3-13
3.8.11	Program Section Attribute ABS Ignored	3-13
3.8.12	Linker ACCVIOs when FP_MODE Literal Missing From Command Line	3-13
3.9	Analyze Utility for OpenVMS	3-13
3.9.1	Unwind Data Display Problem Corrected	3-13
3.9.2	Formatted Symbol Vector Correctly Shown in Data Segment	3-14
3.9.3	Transfer Array Now Formatted in Data Segment	3-14
3.9.4	System Version Array Now Formatted in Dynamic Segment	3-14
3.9.5	Enhancements for the /SEGMENT Qualifier	3-14
3.9.6	Support for Section Escaping Added	3-15
3.10	INSTALL Utility for OpenVMS (Installing Resident Images in S2 Space)	3-15
3.11	Librarian Utility	3-15
3.11.1	Support for Section Escaping	3-15
3.12	InfoServer Utility and FDDI	3-16
3.13	New Qualifier for DCL Command SET PASSWORD	3-16
3.14	INITIALIZE/ERASE=INIT Before Using Media	3-16
3.15	XML-C Product Zip File	3-16
3.16	Using the OpenVMS e-Business and Integration Infrastructure Package	3-16
3.17	Performance Data Collector for OpenVMS (TDC)	3-17
3.18	C++ Compiler	3-17
3.19	Building DCE IDL C++ Applications	3-17
3.20	CDSA Errors During Installation	3-17
3.21	Update to FREEWARE README.TXT File	3-18
3.22	Process/Application Hangs	3-18
3.23	Provisioning	3-18
3.23.1	BL860c System Must Be Powered ON	3-19
3.23.2	InfoServer Name Length	3-19
3.23.3	Provisioning Multiple Servers	3-19
3.23.4	Provisioning from HP SIM Central Management Server	3-19
3.23.5	OpenVMS InfoServer and the Server Blade on the Same Local Area Network	3-19
3.23.6	EFI Firmware	3-19
3.23.7	Management Processor	3-19
3.24	SYSSLDDRIVER Restriction	3-19
3.25	RMS \$PARSE Validation of Directory Files	3-19

3.26	SCACP Error Counter Can Report Retransmit Errors	3-20
3.27	Serial Port Enumeration	3-20
3.28	Old Firmware Cannot Translate Messages Written to the System Event Log	3-22
3.29	AST Delivery Clarification in Programs using POSIX Threads	3-22
3.30	TZ Function in the CRTL	3-23
3.31	HP TCP/IP Services for OpenVMS	3-23
3.31.1	NI_* Flag Values for the getnameinfo() function Improperly Changed	3-23
3.32	Virtual Connect	3-24
3.32.1	Failover and RECNXINTERVAL	3-24
3.32.2	Fibre SCSI	3-24
3.33	Sysgen Parameter CPU_POWER_MGMT Default Has Changed	3-24
3.34	Bootting A Satellite System with Reserved Memory	3-24

4 MSA Utility

4.1	MSA Utility Usage Summary	4-1
4.1.1	Required Privileges	4-1
4.1.2	Restrictions	4-1
4.2	MSA Utility Commands	4-2
	ACCEPT UNIT	4-3
	ADD UNIT	4-4
	DELETE UNIT	4-6
	EXIT	4-7
	FLASH FIRMWARE	4-8
	HELP	4-9
	LOCATE	4-10
	READ FIRMWARE	4-12
	RESET THIS_CONTROLLER	4-13
	RESET OTHER_CONTROLLER	4-14
	SCAN ALL	4-15
	SET CONTROLLER	4-16
	SET GLOBALS	4-17
	SET UNIT	4-19
	SHOW CONNECTIONS	4-23
	SHOW CONTROLLER	4-24
	SHOW DISKS	4-27
	SHOW GLOBALS	4-29
	SHOW PROFILE	4-30
	SHOW OTHER_CONTROLLER	4-31
	SHOW THIS_CONTROLLER	4-32
	SHOW UNIT	4-33
	SHOW UNITS	4-34
	SHOW VERSION	4-35
	START RECOVER	4-36

5 Product Directories

5.1	OpenVMS I64 Operating Environment	5-1
5.1.1	Directories on the OpenVMS I64 Operating Environment DVD	5-1
5.2	Directories on the OpenVMS Freeware CD	5-3
5.3	Open Source Tools CD	5-3
5.3.1	Directories on the Open Source Tools CD.....	5-4
5.4	Product Licensing	5-5
5.5	OpenVMS Version 8.3-1H1 Documentation	5-6

Index

Tables

1-1	Firmware Versions for Entry-Class Integrity Servers	1-8
2-1	Minimum Firmware Required for VGA Console Support by Platform	2-4
3-1	Obsolete and Replacement Parameters	3-1
3-2	Renamed Parameter	3-1
3-3	Obsolete and Replacement Parameters	3-8
5-1	Directory Structure on the OpenVMS I64 Operating Environment DVD	5-2
5-2	OpenVMS Open Source Tools CD	5-5

Preface

About This Guide

This manual describes the new features and release notes for OpenVMS Version 8.3-1H1 for Integrity servers. It also includes information on the MSA utility and updated Linker information.

Intended Audience

This manual is intended for all users of the HP OpenVMS Version 8.3-1H1 operating system for Integrity servers. Read this manual before you install, upgrade, or use OpenVMS Version 8.3-1H1.

Note

Please check the *Cover Letter for HP OpenVMS Version 8.3-1H1 for Integrity Servers* 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:

- Chapter 1 provides an overview of the release, discusses the HP technical support policy, and networking options.
- Chapter 2 describes the new features provided in OpenVMS Version 8.3-1H1 for Integrity servers.
- Chapter 3 contains release notes pertaining to OpenVMS Version 8.3-1H1. Notes are organized by facility or product name when applicable.
- Chapter 4 describes the MSA Utility.
- Chapter 5 provides the directory names and locations of the products on the Operating Environment (OE) DVD and other media included in the OpenVMS Version 8.3-1H1 media kit.

Related Documents

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 OpenVMS Documentation Group, ZKO3-4/Y02 110 Spit Brook Rd. Nashua, NH 03062-2698

How to Order Additional Documentation

For information about how to order additional documentation, see the HP OpenVMS Systems Documentation Ordering page:

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

Conventions

The following conventions may be used in this manual:

Ctrl/ <i>x</i>	A sequence such as Ctrl/ <i>x</i> indicates that you must hold down the key labeled Ctrl while you press another key or a pointing device button.
PF1 <i>x</i>	A sequence such as PF1 <i>x</i> 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: <ul style="list-style-type: none">• Additional optional arguments in a statement have been omitted.• The preceding item or items can be repeated one or more times.• Additional parameters, values, or other information can be entered.
.	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.
<i>italic type</i>	Italic type indicates important information, complete titles of manuals, or variables. Variables include information that varies in system output (Internal error <i>number</i>), in command lines (<i>PRODUCER=name</i>), and in command parameters in text (where <i>dd</i> represents the predefined code for the device type).
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.
UPPERCASE TYPE	Uppercase type indicates a command, the name of a routine, the name of a file, or the abbreviation for a system privilege.
-	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 radixes—binary, octal, or hexadecimal—are explicitly indicated.

Introduction

This chapter introduces the HP OpenVMS Version 8.3–1H1 for Integrity servers product and information that you need to know before installing or upgrading to Version 8.3–1H1.

HP recommends that you read all of the following documents before installing or upgrading to OpenVMS Version 8.3–1H1:

- *HP OpenVMS Version 8.3-1H1 for Integrity Servers New Features and Release Notes* (this manual)
- *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*
- *Cover Letter for HP OpenVMS Version 8.3-1H1 for Integrity Servers*
- *HP OpenVMS Version 8.3 Release Notes*

1.1 Overview of OpenVMS Version 8.3–1H1 for Integrity Servers

OpenVMS Version 8.3–1H1 for Integrity servers (also known as OpenVMS I64) is a follow-on release that replaces OpenVMS Version 8.3 for Integrity servers. OpenVMS Version 8.3–1H1 includes all the capabilities of its predecessor and introduces new features and hardware support.

Major new features include:

- Support for HP BladeSystems c-Class, including new function codes in system services.

With the introduction of HP BladeSystems c-Class, HP has revamped the server blade. This effort encompasses Integrity and ProLiant server blades, as well as storage blades, enclosures, and power and cooling solutions.

Server blades are built with a modular architecture, in which pieces are added to an enclosure to meet the computing needs of the customer.

The enclosure is the metal box that houses the pieces. These enclosures can consist of the following:

- Server blades
- Power supplies
- Fans
- Storage blades
- Onboard administrators
- Interconnect modules

These new server blades are designed to reduce the cost of ownership. You can accomplish this in three primary areas:

Introduction

1.1 Overview of OpenVMS Version 8.3–1H1 for Integrity Servers

- Energy savings
- Configuration and space
- Ease of system management
- Support for the latest Itanium dual-core processors
This release supports the latest Intel Itanium 9100-based systems.
- Support for Virtual Connect
Virtual Connect is a set of interconnect modules and embedded software for HP BladeSystem c-Class enclosures; it simplifies the setup and administration of server connections. HP Virtual Connect includes the HP 1/10Gb Virtual Connect Ethernet Module for HP BladeSystem c-Class, the HP 4Gb Virtual Connect Fibre Channel Module for HP BladeSystem c-Class, and the HP Virtual Connect Manager.

For more information, see

<http://h20000.www2.hp.com/bc/docs/support/SupportManual/c00865618/c00865618.pdf>

- Support for iSCSI Initiator
The iSCSI standard is an emerging Internet Protocol-based storage interconnect standard developed by the Internet Engineering Task Force (IETF). Like Fibre Channel, it is one of a number of standards developed to map the SCSI protocol over a particular transport mechanism.
- System management, including provisioning, providers, and WBEM infrastructure
HP Systems Insight Manager (HP SIM) runs as a management server on ProLiant servers running Windows®. OpenVMS can be managed from HP SIM on any of the supported management servers.
OpenVMS provides software support in the following management areas so that OpenVMS looks and feels similar to server blades running on HP-UX:
 - Provisioning
Provisioning is the automatic configuration (or reconfiguration) of BladeSystems. Using HP SIM with the network services provided by the InfoServer utility and TCP/IP Services for OpenVMS (BOOTP), you can install or upgrade up to eight servers simultaneously. You can also use HP SIM with the vMedia feature provided with your Integrity server firmware, in which case you can install or upgrade OpenVMS on one server at a time.
 - WBEM Infrastructure
This release supports Web Based Enterprise Management (WBEM), whose infrastructure is currently based on OpenPegasus Version 2.5.0.
For more information on WBEM Services, refer to `sys$help:wbem_services*.release_notes`, which is installed with OpenVMS Version 8.3-1H1.
 - WBEM Providers (Instance and Method)
WBEM Services for OpenVMS Version 2.0 is a component of the base operating system in OpenVMS Version 8.3. The WBEM providers comprise the operating system, computer system, process, CPU, memory, fan and power supply, as well as management processor and enclosure. To be managed by HP SIM, your Integrity server requires HP WBEM

1.1 Overview of OpenVMS Version 8.3–1H1 for Integrity Servers

(Web-Based Enterprise Management) Services for OpenVMS and the HP WBEM Providers for OpenVMS. For more information, see *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*.

- Services support (WEBES)
- Support for high-speed USB

This release of OpenVMS includes High Speed, which increases the data transfer rate from the USB-attached DVDs in the 3600/6600 and 2660 systems as well as on future platforms. High Speed allows data to be read from those DVDs more than ten times faster than from USB 1.0/1.1 speeds.

Note

For this release, because the firmware does not yet support High Speed when installing from a DVD, the early stages of the boot are still restricted to USB 1.1 speeds. Until the operating system assumes control of the device, it will run at classic speeds.

Note

For systems using USB keyboards and mouse devices, HP strongly recommends that the keyboards and mouse devices be directly connected to the system and not plugged into a high-speed hub. If a hub is needed for other reasons, HP strongly encourages that it is a FULL-speed hub. Plugging low-speed devices like keyboards and mouse devices can have a major performance impact on the performance of the high-speed hub.

- Support for vMedia

vMedia is the first in a series of virtual-disk capabilities included in the current generation of management-processor hardware. The first platforms that supported vMedia were the rx3600 and rx6600 followed by the rx2660 and finally the c-Class server blades. vMedia consists of two primary components: firmware in MP hardware that emulates a USB DVD, and Java® code that runs on a remote PC that performs I/O to the PC's local CD/DVD drive.

vMedia can be used to do initial system installation on a system without a built-in DVD, or remote installation on a system physically located elsewhere. It also allows the installation of layered products, either remotely or on systems without a built-in DVD.

Note

HP recommends installing vMedia from the ISO file instead of directly installing from the DVD.

- Support for Serial Attached SCSI (SAS) tape blade

OpenVMS Version 8.3-1H1 supports the HP Storageworks Ultrium 488c and Ultrium 920c tape blade for the C-Class Integrity BladeSystem.

Introduction

1.1 Overview of OpenVMS Version 8.3–1H1 for Integrity Servers

Note

SAS tapes must be configured by using the same SYSMAN commands for configuring Fibre Channel tapes (SYSMAN IO FIND_WWID and SYSMAN IO AUTOCONFIGURE). For details, see Section 7.5 of *Guidelines for OpenVMS Cluster Configurations*.

- Support for VGA Console

This release provides a method to boot Integrity systems using a VGA display and a USB keyboard and mouse.

Note

The following restrictions apply to VGA Console support for this release:

- This release does not include XDELTA support or conversational boot support. You must set the graphics as a primary console and one serial line as secondary.
- This release has not been tested on and does not support Keyboard Video Mouse (KVM) switches with VGA.
- VGA Console is supported only for a directly connected monitor and does not support the VGA display available in the Integrated Lights-Out 2 (iLO 2) Web interface.
- When using a VGA console and installing from vMedia or a USB DVD drive with the keyboard plugged into a hub, the keyboard might not be configured. In that event, unplug the hub into which the keyboard is plugged and plug it back in. This should be a rare occurrence and HP plans a correction in the next release.
- Installing from an InfoServer is not supported using VGA Console.
- HP recommends that you do not perform a SYSMAN IO AUTO command on systems with a VGA console. If you have booted either minimum or with the SYSGEN parameter NOAUTOCONFIGURE set and need to configure devices HP recommends that you use the /SELECT qualifier and configure a single device at a time. HP hopes to remove this restriction in a future release.
- If you select "no reboot" when executing the shutdown command procedure, the system still reboots. This restriction will be corrected in a future release. If you do not want the system to reboot, you have the following options:
 - Wait at the console and intercept the system before it executes your default boot option.
 - Use the boot_options command procedure in sys\$manager. With this you can perform the following actions:
 - Change the TIMEOUT to 0 (never time out) or to a large value to provide time to intercept the reboot.
 - Make the EFI Shell the first boot option.

1.1 Overview of OpenVMS Version 8.3–1H1 for Integrity Servers

- Create a boot option that does not actually boot the system.
-

- **Installation Utility**

On all Itanium-based systems, the install utility can install Alpha images. Especially in a mixed-architecture environment, the system manager should be aware of this feature to avoid installing an incorrect image.

Beginning with OpenVMS Version 8.3, the INSTALL LIST command flags such images with the "A*" prefix when displaying the file name of the known image.

- **Smart Array P-800 Lite Support**

OpenVMS Version 8.3-1H1 supports the Smart Array P-800 Lite P800 (AD335A) 16-port Serial Attached SCSI (SAS) RAID controller PCI Express.

- **4-GB Fibre Channel PCIe Adapter Support**

OpenVMS Version 8.3-1H1 supports 1-port 4-Gb Fibre Channel Adapter (AD299A) and 2-port 4-GB Fibre Channel Adapter (AD355A) PCI Express based on the Emulex LPe11002 chip.

- **Graceful Shutdown**

This capability allows you to initiate a "graceful shutdown" from the MP console. OpenVMS then institutes a shutdown, which ends with Primary halted with code HWRPB_HALT\$K_POWEROFF. Depending on the platform, this leaves the MP running or actually powers off the system.

- **Edgeport/8 and EdgePort/416 Multiplexer Support**

The Edgeport/8 and Edgeport/416 (8- and 16-line serial multiplexers) from InsideOut Networks (a Digi International Company) have been replaced to meet RoHS (Reduction of Hazardous Substances) directive of the European Union.

- **MSA utility**

MSA utility is an OpenVMS system management tool used for configuring and managing the following controllers:

- HP StorageWorks SmartArray Family of Storage Host Bus Adapters (5300 series, 6400 series, P400, and P800 Lite)
- HP StorageWorks Modular Smart Array Family of Storage Controllers (MSA1000 and MSA1500)

These controllers connect to HP StorageWorks Modular Smart Array storage array systems and to Integrity server internal SCSI or SAS drives.

Note

You must have the following privileges to use this utility:

NETMBX, TMPMBX, SYSPRV, DIAGNOSE, PHY_IO

Introduction

1.1 Overview of OpenVMS Version 8.3–1H1 for Integrity Servers

The new devices look identical to the old devices and have the same order numbers; however, the old devices use an ION930 chip while the new devices use the TI (Texas Instruments) TUSB5152 chip. The TUSB5152 is incompatible with the existing OpenVMS device driver. A new device driver provides support for the new devices in a relatively transparent fashion.

The new terminal devices will continue to be named TXD and will provide the same performance and capabilities of the old Edgeport devices. New and old Edgeport devices can be mixed in the same system; however, users of the Edgeport/416 might need to check the naming order of TXD devices and rename them if they do not meet their needs.

1.2 HP Software Technical Support Policy

Unless otherwise agreed to by Hewlett-Packard Company (HP), HP provides HP Services support for only the current and immediately preceding versions of HP software, and only when the software is used with hardware that is included in HP-specified configurations. A version is defined as a release of a software product that contains new features, enhancements, or maintenance updates.

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 I64, as well as for Alpha and OpenVMS VAX, is kept up to date on this website:

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

The following OpenVMS core products are supported at the same level (SS or PVS) and duration as the operating system version on which they ship:

- 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 RMS Journaling for OpenVMS
- HP TCP/IP Services for OpenVMS
- HP Volume Shadowing for OpenVMS

These products must be listed as line items on an HP software support agreement in order for you to receive software support on these products. The DECnet Phase IV product requires a Prior Version Support line item.

1.3 General Application Compatibility Statement

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

1.4 Obtaining Patch Kits

Patch kits, also known as remedial kits, for HP products are available on line 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://www1.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/i64/V8.3-1H1

1.5 Networking Options

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

During the main installation procedure for OpenVMS Version 8.3-1H1, you have the option of installing the following supported HP networking software:

- Either HP DECnet-Plus Version 8.3-1H1 for OpenVMS or HP DECnet Phase IV for OpenVMS. (Note that these 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.

Standard support for DECnet Phase IV is provided to customers. For more information about the Prior Version Support service, see Section 1.2.

- HP TCP/IP Services for OpenVMS Version 5.6ECO2

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. For more information about running DECnet over TCP/IP (RFC 1859) and OSI over TCP/IP (RFC 1006), see the *DECnet-Plus for OpenVMS Management Guide*.

Alternatively, after you install OpenVMS, you can install your choice of another third-party networking product that runs on OpenVMS Version 8.3-1H1.

For information about how to configure and manage your HP networking software after installation, see the TCP/IP, DECnet-Plus, or DECnet documentation. The manuals are available in online format on the OpenVMS Documentation CD, the Online Documentation Library CD, and the OpenVMS Documentation website:

<http://h71000.www7.hp.com/doc/index.html>

To order printed manuals from HP, see the Preface.

1.6 System Event Log on Integrity Servers

HP Integrity servers maintain a System Event Log (SEL) within system console storage, and OpenVMS I64 automatically transfers the contents of the SEL into the OpenVMS error log. If you are operating from the console during a successful boot operation, you might see a message indicating that the Baseboard Management Controller (BMC) SEL is full. You can safely continue when the

Introduction

1.6 System Event Log on Integrity Servers

BMC SEL is full by following the prompts; OpenVMS automatically processes and clears the contents of the SEL.

1.7 Firmware for Integrity Servers

OpenVMS Version 8.3–1H1 was tested with the latest firmware for each of the supported Integrity servers.

For the entry-class Integrity servers, HP recommends that use the most current system firmware. For information about updating the system firmware for entry-class Integrity servers, see the *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*. (For rx7620, rx8620, and Superdome servers, call HP Customer Support to update your firmware.)

Table 1–1 lists the recommended firmware versions for entry-class Integrity servers:

Table 1–1 Firmware Versions for Entry-Class Integrity Servers

System	System Firmware	BMC Firmware	MP Firmware	DHCP Firmware
rx1600	4.27	4.01	E.03.30	N/A
rx1620	4.27	4.01	E.03.30	N/A
rx2600	2.31	1.53	E.03.30	N/A
rx2620	4.27	4.04	E.03.30	N/A
rx4640	4.28	4.04	E.03.30	1.10
rx2660*	1.05	5.06	F.01.58	N/A
rx3600*	2.03	5.14	F.01.58	N/A
rx6600*	2.03	5.14	F.01.58	N/A

*If you have Intel Itanium 9100 processors on your rx2660, rx3600, or rx660, you need firmware that is at least one version greater than the ones listed here.

For cell-based servers, you must access the MP Command Menu and issue the `sysrev` command to list the MP firmware revision level. The `sysrev` command is available on all HP Integrity servers that have an MP. Note the `EFI info fw` command does not display the Management Processor (MP) firmware version on cell-based Integrity servers.

To check firmware version information on an entry-class Integrity server that does not have the MP, enter the `info fw` command at the EFI prompt. Note the following example:

```
Shell> info fw
```

```
FIRMWARE INFORMATION
```

```
Firmware Revision: 2.13 [4412]
```



```
PAL_A Revision: 7.31/5.37
```

```
PAL_B Revision: 5.65
```

```
HI Revision: 1.02
```

```
SAL Spec Revision: 3.01
```

```
SAL_A Revision: 2.00
```

```
SAL_B Revision: 2.13
```

```

EFI Spec Revision: 1.10
EFI Intel Drop Revision: 14.61
EFI Build Revision: 2.10

POSSE Revision: 0.10

ACPI Revision: 7.00

BMC Revision: 2.35                ❷
IPMI Revision: 1.00
SMBIOS Revision: 2.3.2a
Management Processor Revision: E.02.29  ❸

```

- ❶ The system firmware revision is 2.13.
- ❷ The BMC firmware revision is 2.35.
- ❸ The MP firmware revision is E.02.29.

The HP Integrity rx4640 server contains Dual Hot Plug Controller (DHPC) hardware with upgradable firmware. To check the current version of your DHPC firmware, enter the EFI command `info chiprev`, as shown in the following example. The hot-plug controller version will be displayed. A display of 0100 indicates version 1.0; a display of 0110 indicates version 1.1.

```
Shell> info chiprev
```

```
CHIP REVISION INFORMATION
```

Chip Type	Logical ID	Device ID	Chip Revision
Memory Controller	0	122b	0023
Root Bridge	0	1229	0023
Host Bridge	0000	122e	0032
Host Bridge	0001	122e	0032
Host Bridge	0002	122e	0032
Host Bridge	0004	122e	0032
HotPlug Controller	0	0	0110
Host Bridge	0005	122e	0032
HotPlug Controller	0	0	0110
Host Bridge	0006	122e	0032
Other Bridge	0	0	0002
Other Bridge	0	0	0008
Baseboard MC	0	0	0235

For instructions on how to access and use EFI, see the *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*. For more information, refer to the hardware documentation that is provided with your server.

For instructions on upgrading your firmware for your entry-class Integrity servers, refer to the *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*. To upgrade firmware for the rx7620, rx8620, or Superdome, contact HP Customer Support.

1.8 Release Notes on Booting the System

The following release notes pertain to booting the OpenVMS I64 system.

Introduction

1.8 Release Notes on Booting the System

1.8.1 Booting from the Installation DVD

On I64 systems with the minimum amount of supported memory, the following message appears when booting from the installation DVD:

```
***** XFC-W-MemmgmtInit Misconfigure Detected *****
XFC-E-MemMisconfigure MPW_HILIM + FREEGOAL > Physical Memory and no reserved memory for XFC
XFC-I-RECONFIG Setting MPW$GL_HILIM to no more than 25% of physical memory XFC-I-RECONFIG
Setting FREEGOAL to no more than 10% of physical memory
***** XFC-W-MemMisconfigure AUTOGEN should be run to correct configuration *****
***** XFC-I-MemmgmtInit Bootstrap continuing *****
```

The message means that the system cache (XFC) initialization has successfully adjusted the SYSGEN parameters MPW_HILIM and FREEGOAL to allow caching to be effective during the installation. You can continue with the installation.

1.8.2 Setting Up I64 Systems to Reboot

An OpenVMS I64 system does not reboot automatically unless you have it set up to do so either by using EFI or by using the OpenVMS I64 Boot Manager utility.

For information about how to set up your I64 system to automatically reboot, refer to the *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*.

1.8.3 Booting with a Common Cluster System Disk

For configuring additional nodes to boot with a common cluster disk, refer to the CLUSTER_CONFIG_LAN utility described in the *HP OpenVMS System Manager's Manual, Volume 1: Essentials*.

For additional information about the I64 Boot Manager Boot Options Management Utility, see the *HP OpenVMS System Manager's Manual, Volume 1: Essentials*.

1.8.4 Booting from a Fibre Channel Storage Device

Many customers prefer to boot from a Fibre Channel (FC) storage device because of its speed and because it can serve as a common cluster system disk in a SAN. Booting from an FC storage device on OpenVMS I64 systems is significantly different from booting from an FC storage device on OpenVMS Alpha systems.

For instructions on how to configure and boot from an FC device on OpenVMS I64 systems, see the Fibre Channel appendix of *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*.

1.8.5 OpenVMS I64 Boot Manager Utility: Adding Multipath Fibre Channel Disk Devices

The OpenVMS Boot Manager utility, BOOT_OPTIONS.COM, is used to specify a list of boot and dump devices in a SAN. When you add a multipath Fibre Channel disk device to the list, all paths to the device found in the SAN, including redundant paths, are listed.

1.8.6 Fibre Channel Boot Disk: Simplified Setup Process

For OpenVMS I64 Version 8.3-1H1, the process of setting up a Fibre Channel boot device requires the use of the OpenVMS I64 Boot Manager, `BOOT_OPTIONS.COM`, to specify values to the EFI Boot Manager. This process is automated by the OpenVMS I64 Version 8.2-1 installation process, although the manual process is still available for those cases when it might be needed.

The OpenVMS I64 Version 8.3-1H1 installation process displays the name of a Fibre Channel disk as a boot device and prompts you to add the Boot Option. HP recommends that you accept this default. Alternatively, you can run the OpenVMS I64 Boot Manager after the installation or upgrade completes, as described in the *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*.

Note

If your system is a member of the rx1600, rx2600, or rx4600 family of servers and a Fibre Channel boot device is not listed in the EFI boot menu, you might experience a delay in the EFI initialization because the entire SAN is scanned.

Depending on the size of the SAN, this delay can range from several seconds to several minutes. Cell-based systems (the rx7620, rx8620, and Superdome families of servers) are not affected by this delay. This delay might occur when booting OpenVMS from the installation DVD for the first time on any OpenVMS I64 system.

For information on booting from a Fibre Channel boot device and updating the Fibre Channel adapter firmware, see the *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*.

1.9 HP DECwindows Motif

The following DECwindows Motif release notes are of interest to OpenVMS I64 users.

1.9.1 Connect Peripheral Devices Prior to Server Startup

To properly configure your system as a DECwindows X display server, you must have all the following peripheral components connected prior to startup:

- Monitor
- USB mouse
- USB keyboard

Otherwise, the server system might not complete the device initialization process correctly. For example, starting up a server system without input devices (mouse and keyboard) results in a blank screen, and DECwindows does not start.

To correct this problem, connect all peripherals. Then restart the server, or reboot the system. Note that only one keyboard and one mouse are supported. For DECwindows to detect and use these peripherals, they must show as `KBD0` and `MOU0` when doing `SHOW DEVICE KBD` and `SHOW DEVICE MOU` commands.

Introduction

1.9 HP DECwindows Motif

1.9.2 Countdown Messages Displayed During Startup

When running DECwindows Motif in client-only mode (with no server configured, or with no mouse or keyboard connected), messages similar to the following might be displayed during startup:

```
Waiting for mouse...
Waiting for keyboard...
```

These messages indicate that device polling is underway and are informational only. They will disappear when the 15-second countdown is complete. This typically occurs on servers that incorporate built-in graphics as part of a Server Management Option, and no keyboard and mouse are connected.

To prevent the messages from displaying, and the 15-second delay, connect the input devices (USB mouse and USB keyboard) to the system prior to startup.

If you do not intend to use local graphics on the system, you can define the logical name in SYSTARTUP_VMS.COM as follows:

```
$ DEFINE/SYSTEM DECW$IGNORE_WORKSTATION TRUE
```

This prevents the DECwindows startup from attempting to start local graphics operation.

1.9.3 Optional Graphics

The ATI Radeon 7500 PCI option (HP part number AB551A) is supported on the entry-class Integrity servers for 2D multi-head and 3D operation. Refer to the Installation Guide for this device or Quickspecs for information about configuration and use.

1.9.4 Keyboard Support

The OpenVMS USB keyboard (HP part number AB552A) is supported on all Integrity systems supporting DECwindows graphics. The keyboard is packaged with a three-button thumbwheel mouse.

1.9.5 Firmware Update and Keyboards

If you update the firmware on your system before upgrading to OpenVMS Version 8.3-1H1, the system might think it has more than one keyboard. If this happens, DECWindows might not start or might not have a usable keyboard. This applies to the rx2660, rx3600, rx6600, and BL860c systems.

To correct this problem, install OpenVMS Version 8.3-1H1, then delete all copies of SYSSSPECIFIC:[SYSEXE]USB\$UCM_DEVICES.DAT.

Note

Be sure not to delete SYSSCOMMON:[SYSEXE]USB\$UCM_DEVICES.DAT. The newer firmware causes the keyboard in the presently unsupported vKVM hardware to become visible to OpenVMS. OpenVMS Version 8.3-1H1 has code to ignore the latent keyboard, a feature that OpenVMS Version 8.3 did not have.

OpenVMS Version 8.3–1H1 New Features

This chapter describes the new features provided in OpenVMS I64 Version 8.3–1H1.

2.1 ISV Applications and Binary Compatibility in HP OpenVMS Version 8.3-1H1

With the HP OpenVMS hardware release (Version 8.3-1H1), HP wants to assure ISVs and Layered Product developers that existing application binaries that work today on OpenVMS Version 8.3 will continue to run unchanged on OpenVMS Version 8.3-1H1. OpenVMS Version 8.3-1H1 is engineered to maintain binary compatibility with OpenVMS version 8.3. The thousands of applications written for OpenVMS version 8.3 will continue to work unchanged on OpenVMS version 8.3-1H1, through binary compatibility ensured by quality operating system engineering and lab testing.

Software recompilation is not required, nor will ISVs or Layered Product developers have to retest or requalify their products against Version 8.3-1H1 unless they wish to do so. If you have qualified on OpenVMS Version 8.3, you are already qualified on OpenVMS Version 8.3-1H1.

If ISVs or Layered Product developers uncover problems on OpenVMS Version 8.3-1H1 with products that work on OpenVMS Version 8.3, notify HP immediately. HP is committed to resolving any compatibility issues found with OpenVMS Version 8.3-1H1. .

2.2 New Item Code `DVI$_ADAPTER_IDENT`

On Alpha and I64 systems, this returns (as a string) the description of an adapter as defined in either `SYSSYSTEM:SYS$CONFIG.DAT` or `SYSSYSTEM:SYS$USER_CONFIG.DAT`. Note that this service does not read either of those files; those files are read into memory in response to the `SYSMAN IO REBUILD` command, which typically happens while a system is booting.

HP recommends a buffer size of 255 bytes to hold the identification string.

2.3 New item code `DVI$_MOUNTCNT_CLUSTER`

On Alpha and I64 systems, this item code returns (as a longword) the number of systems in a cluster that have a device mounted. Note that this new item code is not a direct replacement for the existing item code `DVI$_MOUNTCNT`. That item code returns the number of mounters for any device on the local system. The `/SHARE` qualifier to the `MOUNT` command can allow for more than one mounter.

OpenVMS Version 8.3–1H1 New Features

2.4 HP Smart Array P800 Controller (AD335A) Support

2.4 HP Smart Array P800 Controller (AD335A) Support

OpenVMS V8.3-1H1 supports the Smart Array P800 16-port serial attached SCSI (SAS) controller with PCI-Express (PCIe). The HP StorageWorks 60 Modular Storage Array (MSA60) and 70 Modular Storage Array (MSA70) SAS storage enclosure can be connected to the external ports of this controller.

For additional information see the following web sites:

<http://www.hp.com/products1/serverconnectivity/storagesnf2/sas/index.html>

http://h18004.www1.hp.com/storage/disk_storage/msa_diskarrays/drive_enclosures/index.html

2.5 HP 4GB Fibre Channel Adapters (AD299A/AD355A) Support

OpenVMS V8.3-1H1 supports the 2-port 4Gb Fibre Channel Adapter and 1-port 4Gb Fibre Channel controller with PCI-Express (PCIe).

For additional information, see the following web site:

<http://www.hp.com/products1/serverconnectivity/storagesnf2/4gbfibre/index.html>

2.6 HP StorageWorks Ultrium 448c Tape Blade

OpenVMS V8.3-1H1 supports the Ultrium 448c half-height tape blade providing an integrated data protection solution for HP BladeSystem c-Class enclosures.

For additional information, see the following web site:

<http://h18004.www1.hp.com/products/servers/storageworks/c-class/ultrium-448c/index.html>

2.7 Storage and Network I/O Controllers supported after the Initial OpenVMS V8.3 Release

The following list of storage and network I/O controllers were released after OpenVMS V8.3 shipped. V8.3 support was introduced via patch kits; now all of these devices are fully qualified and supported by OpenVMS V8.3-1H1.

- HP Smart Array P400
- HP StorageWorks SB40c Storage Blade
- 2-port 4Gb FC PCIe (AD300A)
- 2-port 4Gb FC Mezzanine (403619-B21)
- 1-port 4Gb FC & 1 port GigE HBA PCI-X combo (AD193A)
- 2-port 4Gb FC & 2 port GigE HBA PCI-X combo (AD194A)
- 1-Port 1000Base-R PCI-X (AD331A)
- 2-Port 1000Base-R PCI-X (AD332A)

2.8 Graphics Console Support for Selected HP Integrity Server Platforms

OpenVMS V8.3-1H1 provides support for a graphics console on selected HP Integrity platforms. This functionality allows the user to connect a monitor, keyboard, and mouse directly to connectors provided for that purpose on the bulkhead of the system. Previous versions of OpenVMS allowed console

OpenVMS Version 8.3–1H1 New Features

2.8 Graphics Console Support for Selected HP Integrity Server Platforms

connections only via a serial terminal connected to a serial port on the bulkhead or through a terminal emulator through the network interface.

Additionally, the user may select a graphics option card in a PCI slot to be the graphics console. Some HP Integrity server systems might impose limitations on this capability and on the number of graphics PCI cards for multiple-head graphics capabilities. For platform-specific, option-card limitations on graphics console and multiple graphics heads, see your HP Integrity server documents.

Note

Some platforms do not have an embedded graphics device. For those platforms, a graphics option card is required to affect a graphics console.

2.8.1 Graphics Console Support Platforms

Some older platforms, such as the rx2600, do not support the graphics console, because they lack the firmware capabilities required for OpenVMS graphics console support.

Currently, the OpenVMS graphics console supports the following HP Integrity server platforms:

- rx1600
- rx1620
- rx2620
- rx4640
- rx3600
- rx6600
- rx2660
- BL860c Server Blade

2.8.2 Supported PCI Graphics Option Cards

At this time, one PCI Graphics option card is supported for the listed HP Integrity servers (HP Product Number AB551A). This is the ATI Radeon PCI graphics card.

2.8.3 Graphics Console Limitations

The Graphics Console currently has the following limitations:

- Conversational Boot for Graphics Console is not supported on IA64 platforms. Setting bit 0 of the boot flags, `BootFlag(0)`, to enable conversational boot elicits an informational message from the Boot Loader, which clears the Conversational Boot flag automatically before proceeding.
- XDELTA Boot for Graphics Console is not supported on IA64 platforms. Setting bit one of the boot flags, `BootFlag(1)`, elicits an informational message from the Boot Loader. The Boot Loader automatically clears this bit in the Boot Flags variable before proceeding.

OpenVMS Version 8.3–1H1 New Features

2.8 Graphics Console Support for Selected HP Integrity Server Platforms

2.8.4 Minimum Firmware Required for Graphics Console Support

Table 2–1 lists the minimum firmware revisions for Graphics Console support for the referenced platforms as of the date of this document.

Table 2–1 Minimum Firmware Required for VGA Console Support by Platform

HP Integrity Server	System Firmware	Baseboard Management Controller	Management Processor
rx1600	4.13	4.01	E03.30
rx2620	4.21	4.03	E.03.30
rx4640	4.21	4.05	E.03.30
rx3600, rx6600	2.03	5.14	F.01.58
rx2660	1.05	5.02	F.02.58
BL860c	1.01	74.52	T2.02

System Firmware, Baseboard Management Controller, and Management Processor firmware are updated with a single firmware image.

You can determine the firmware revision on your system in one of the following ways:

- From the EFI Shell> prompt
- From the MP:CM> prompt
- From the cli> prompt
- From the EFI Boot Manager Menu (Refer to *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*.)

All systems will have an EFI interface and one or both Management Port (MP) and BMC port (cli) interfaces. The examples that follow were taken from an rx2620.

2.8.4.1 Obtaining Firmware Information from the EFI Shell

To obtain firmware information from the EFI shell, enter `info fw` command at the Shell> prompt:

```
Shell> info fw
-
FIRMWARE INFORMATION
-
System Firmware Revision: 4.03 [4624]
BMC Revision: 4.01
Management Processor Revision: E.03.28
```

2.8.4.2 Obtaining Firmware Information from the MP Prompt

Note

Verify that your system has the Management Port option. For information about your server, see your system's documentation.

OpenVMS Version 8.3–1H1 New Features

2.8 Graphics Console Support for Selected HP Integrity Server Platforms

Whether connecting to the Management Port remotely over the network or through the serial port's DE9 connector provided for the MP on the rear bulkhead, you must first enter the user name and password to access the MP prompt. If you are already at the EFI Shell> prompt, press `ctrl/b`.

To obtain firmware information from the MP prompt, enter the `cm` command at the `MP>` prompt:

```
MP> cm
-
          (Use Ctrl-B to return to MP main menu.)
-
MP:CM>
```

At the `MP:CM>` prompt, enter the `sysrev` command:

```
MP:CM> sysrev
-
SYSREV
-
Current firmware revisions
-
MP FW      : E.03.28
BMC FW     : 04.01
EFI FW     : 05.46
System FW  : 04.03
```

2.8.4.3 Obtaining Firmware Information from the cli Prompt

The BMC serial port that provides the `cli>` prompt is not available on all systems. For information about your system, see your system's documentation. Connect a serial terminal to the DE9 connector provided for the BMC port on the rear bulkhead. To connect to the `cli>` interface, enter the `esc-()` command. From the `cli>` prompt, enter the `sr>` command:

```
cli> sr
Current firmware revisions:
-
MP FW      E.03.28
BMC_FW     04.01
EFI FW     05.46
System FW  04.03
```

2.8.4.4 Obtaining Firmware Information from the EFI Boot Manager

To determine the firmware version, see the System Overview panel in the EFI Boot Manager.

2.8.5 Platform-Specific Connection Details

Note

To implement multiple head graphics, you must rename the file `DECW$PRIVATE_SERVER_SETUP.TEMPLATE` to `DECW$PRIVATE_SERVER_SETUP.COM`.

Verify that the mouse and keyboard are plugged into USB connectors provided on the bulkheads, according to the following instructions for each platform.

Do not use the DVI connector for multiple-head operation.

Use the analog DB15 connector provided on the option graphics card.

OpenVMS Version 8.3–1H1 New Features

2.8 Graphics Console Support for Selected HP Integrity Server Platforms

2.8.5.1 HP Integrity rx1600 and rx1620 Servers

Note

The embedded graphics device is optional on these systems. For more information, see your system's documentation.

These systems might not have the optional embedded graphics device. If this is the case, insert a graphics option card into the PCI option slot in the system to provide a graphics console. Although two PCI options slots are available, only one is large enough to accommodate a graphics card. Thus, multiple-head support is not available without the embedded graphics device option.

Connect the monitor to the Graphics DB15 graphics connector and connect the mouse and keyboard to the USB bulkhead connectors.

2.8.5.2 HP Integrity rx2620 and rx4640 Servers

Note

The embedded graphics device is optional on these systems. For more information, see your system's documentation.

The DB15 analog connector to the embedded graphics controller and USB connectors for the mouse and keyboard are provided on the rear bulkhead.

To obtain multihead operation, insert option graphics cards into available PCI slots.

Do not use DVI connectors on the option graphics cards; instead, use the DB15 analog connectors.

2.8.5.3 HP Integrity rx3600 and HP Integrity rx6600 Servers

Note

The embedded graphics device is optional on these systems. Check your system documentation.

The DB15 analog connector to the embedded graphics controller and USB connectors for the mouse and keyboard are provided on the rear bulkhead. Additionally, you can use the front bulkhead USB connector that can be used for the mouse or keyboard.

To obtain multiple-head operation, insert option graphics cards into available PCI slots.

Do not use DVI connectors on the option graphics cards; instead, use the DB15 analog connectors.

When using the internal DVD drive to install OpenVMS from the graphics console, either the keyboard must be plugged into the USB connector on the front panel and the mouse must be plugged into any USB on the rear bulkhead, or you must connect the mouse and keyboard to the two rear bulkhead USB connectors.

2.8 Graphics Console Support for Selected HP Integrity Server Platforms

2.8.5.4 HP Integrity rx2660 Servers

Note

The embedded graphics device is optional on these systems. For more information, see your system's documentation.

The DB15 analog connector to the embedded graphics controller and USB connectors for the mouse and keyboard are provided on the rear bulkhead. Additionally, there is a front bulkhead USB connector that can be used for the mouse or keyboard.

To obtain multiple-head operation, insert option graphics cards into available PCI slots.

Do not use DVI connectors on the option graphics cards; instead, use the DB15 analog connectors.

When using the internal DVD drive to install OpenVMS from the graphics console, you can use the DB15 and USB connectors on the front panel to connect the monitor and keyboard, while using the rear bulkhead USB connector for the mouse.

If you prefer to use the rear bulkhead connectors for the monitor, keyboard, and mouse, you must connect the mouse and keyboard to the two rear bulkhead USB connectors.

2.8.5.5 HP Integrity BL860c Server Blade

The BL860c system comes with an SUV dongle that provides a DB15 connector for the embedded Graphics Console Device and two USB connectors for the keyboard and mouse.

If you want to install OpenVMS with an external USB DVD drive, you can connect the drive to either of the USB connectors and the keyboard to the other. The mouse is not needed for the installation procedure, because the procedure presents a character-based menu.

To obtain multihead operation, insert option graphics cards into available PCI slots.

Do not use DVI connectors on the option graphics cards; instead, use the DB15 analog connectors.

2.9 OpenVMS Internationalization Data Kit (VMSI18N)

The OpenVMS Version 8.3-1H1 internationalization data kit has the following new iconv converters:

- DECKANJI2000_UCS-2.ICONV
- DECKANJI2000_UCS-4.ICONV
- DECKANJI2000_UTF-8.ICONV
- UCS-2_DECKANJI2000.ICONV
- UCS-4_DECKANJI2000.ICONV
- UTF-8_DECKANJI2000.ICONV

OpenVMS Version 8.3–1H1 New Features

2.10 LDAP External Authentication

2.10 LDAP External Authentication

LDAP External Authentication provides the following enhancements.

2.10.1 Enhanced LOGINOUT.EXE and SETP0.EXE Images for LDAP External Authentication

This release provides optional `LOGINOUT.EXE` and `SETP0.EXE` (`SET PASSWORD`) images that utilize the `SYS$ACM` system service for user authentication and password changes.

When these images are used, login and password change requests are sent to the `SYS$ACM` service and handled by the `ACME_SERVER` process's authentication agents.

An OpenVMS authentication agent is configured by default to service standard OpenVMS login and password-change requests. In addition, you can install an LDAP authentication agent that services login and password-change requests using an LDAP version 3 directory server.

2.10.2 Active Directory Support

LDAP external authentication includes additional authentication and password change support for Active Directory's LDAP operations.

For more information, see `SYS$HELP:ACME_DEV_README.TXT`.

2.11 Linker New Features

The following sections describe enhancements to the OpenVMS Linker for this release.

2.11.1 Support for Larger I/O Segments

In OpenVMS Version 8.3, the linker limited the I/O segment to 65535 pagelets, although the I64 image (dynamic segment) allowed you to specify greater values. To enable use of larger buffers, the upper limit of the I/O segment has been lifted. The OpenVMS I64 Linker now allows up to 512MB (half of P1 space) for the I/O segment.

2.11.2 Optimizing Linker Generated Code

For OpenVMS I64, the linker must generate code stubs so that compiler-generated PC relative branches can be used for branching to further offsets and for calling nonlocal routines. The OpenVMS Version 8.3 linker creates a code stub per branch/call. In other words, for multiple calls to `printf`, the linker creates multiple code stubs.

For OpenVMS Version 8.3-1H1, the linker was enhanced to re-use as much of its already-created code stubs as possible. For multiple calls to `printf`, the changed linker creates a single code stub.

In addition, the OpenVMS Version 8.3-1H1 linker places all code stubs at the end of a code segment. As a result, unused code stubs can be trimmed back.

The Program Section Synopsis of the linker map shows this enhancement. Most often, code stubs appear in a linker-created PSECT: `"$LINKER C$n"`, `n=0,...`. In case of several code segments, several of these sections appear. As with the previous linker, the Version 8.3-1H1 linker provides additional Linker code contributions for large code segments.

The new linker-generated code PSECTs are not user changeable.

2.11.3 Signature Segment Flagged in Map File

If the linker creates a signature segment to allow interoperability with translated images, the segment is flagged in the Image Segment Synopsis of the linker map.

2.11.4 Unwind Segments Follow Code Segments

If a user instructs the linker to move code into P2 space (/SEGMENT_ATTRIBUTE=CODE=P2), the Version 8.3.1–1H1 linker also moves the corresponding unwind segments into P2 space. Usually, code is moved to the P2 region to free up virtual address space in the P0 region. With this enhancement, additional address space can be reclaimed from P0 space.

2.12 System Dump Analyzer (SDA) Utility

The following sections list new SDA features provided with the OpenVMS I64 Version 8.3–1H1 operating system. For descriptions of the SDA commands and performance information, see *VMS System Dump Analyzer Utility Manual*.

2.12.1 Changes to CLUE Utility

The CONFIG/CPU parameter of the CLUE command has the following new description:

Displays only the part of the system configuration that contains information about the CPUs.

2.12.2 CLUE Utility: New Qualifier

The /MEMORY parameter of the CLUE command has been added with the following new description:

Displays only the part of the system configuration that contains information about the layout of physical memory.

System information is now printed on the "first page" of a plain CLUE CONFIG command.

The output under CLUE CONFIG/CPU now does not contain the system or memory data. On Integrity systems, the layout of the actual CPU-specific information has also changed.

2.12.3 SHOW_PFN_DATA Command: New Qualifier

A new /SUMMARY[=PROCESS] qualifier has been added to the SDA command SHOW PFN_DATA. The SHOW PFN_DATA/SUMMARY command produces a single-page summary of all pages on the system. Adding the PROCESS keyword causes SDA to generate two additional lists: one for all processes in the system showing the breakdown of its active pages; and one for all processes that have non-active pages (for example, pages on the modified list).

2.12.4 SHOW SUMMARY Command: New Qualifier

Two new qualifiers, /PAGES and /TOTALS, have been added to the SDA SHOW SUMMARY command. SHOW SUMMARY/PAGES causes an additional line to be output, for each process, that gives a breakdown of the process-private and global pages in the process's working set. SHOW SUMMARY/TOTALS will output a summary of process and global pages for resident and nonresident processes, followed by a summary of the processes (or threads if /THREADS was

OpenVMS Version 8.3–1H1 New Features

2.12 System Dump Analyzer (SDA) Utility

also specified) in each scheduling state. /PAGES and/TOTALS may be specified together.

2.12.5 SET AND SHOW CPU Commands

- The SET and SHOW CPU commands now provide limited functionality about running systems: display CPU database address and set CPUDB symbol to the database. The address is also displayed and set when SET CPU is used in a dump.
- The /FIRST, /NEXT, and /PRIMARY qualifiers have been added to both commands:
 - The /FIRST qualifier sets or shows the lowest-numbered CPU (not necessarily the primary).
 - The /NEXT qualifier sets or shows the next CPU that exists. It allows users to walk through all CPUs in the system and displays a BADCPU error when users go beyond the last CPU. This is similar to the SET PROCESS/NEXT command, as shown in the following example:

```
$ CREATE A.COM
FORMAT CPUDB
SET CPU/NEXT
^Z
$ ANALYZE/CRASH
...
SDA> SET CPU /FIRST
SDA> @A
SDA> REPEAT/UNTIL=BADCPU
```

See also the REPEAT command.

- The /NOLOG qualifier has been added to the SET CPU command. This qualifier inhibits messages generated by the command.

2.13 Time Zones Added

OpenVMS Version 8.3-1H1 provides 552 time zones based on the time-zone public database named tzdata2006n. The following eight time zones have been added in OpenVMS Version 8.3-1H1:

- America/Atikokan
- America/Blanc-Sablon
- America/North_Dakota/New_Salem
- Europe/Guernsey
- Europe/Isle_of_Man
- Europe/Jersey
- Europe/Podgorica
- Europe/Volgograd

For a list of the time-zone names, see the *HP OpenVMS System Manager's Manual*.

OpenVMS Version 8.3–1H1 Release Notes

This chapter contains the OpenVMS Version 8.3–1H1 release notes. Because most of the OpenVMS Version 8.3 release notes still apply, read the release notes that apply to OpenVMS I64 (that is, all notes not marked "Alpha only") in the *HP OpenVMS Version 8.3 Release Notes* manual as well as the notes in this chapter before installing or upgrading your system.

Additionally, for further information on booting the system, see Section 1.8.

3.1 OpenVMS Version 8.3-1H1 System Parameters

The following sections describe obsolete, renamed, and new system parameters for OpenVMS 8.3–1H1.

3.1.1 Obsolete Parameters

Table 3–1 lists system parameters that are obsolete in OpenVMS Version 8.3-1H1. The following sections contain information on the replacement parameters.

Table 3–1 Obsolete and Replacement Parameters

Obsolete in OpenVMS V8.3	OpenVMS V8.3–1H1 Replacement
SMP_CPUS	SMP_CPU_BITMAP
SMP_CPUSH	-----
IO_PREFER_CPUS	IO_PRCPU_BITMAP

3.1.2 Renamed Parameters

The following system parameter has been renamed and supplemented in OpenVMS Version 8.3–1H1. Table 3–2 contains additional information about the new parameters.

Table 3–2 Renamed Parameter

Old Name in OpenVMS V8.3	New Names in OpenVMS V8.3–1H1
SCH_CTLFLAGS	SCHED_FLAGS
	SCH_SOFT_OFFLD
	SCH_HARD_OFFLD

OpenVMS Version 8.3–1H1 Release Notes

3.1 OpenVMS Version 8.3-1H1 System Parameters

3.1.3 New Parameter

This release also contains the new `GH_RES_CODE_S2` parameter, which specifies the size in pages of the resident 64-bit S2 space resident image code granularity hint region.

Only images linked with the `/SEGMENT=CODE=P2` qualifier can have code placed in this region. See the *HP OpenVMS Linker Utility Manual* and the `INSTALL` utility in the *HP OpenVMS System Manager's Manual* for more information.

`GH_RES_CODE` has the `AUTOGEN` and `FEEDBACK` attributes.

3.1.4 Failure of `AUDIT_SERVER` to Initiate During Boot

During boot, if the `AUDIT_SERVER` fails to initiate for any reason, the startup process enters a retry loop that attempts to restart the `AUDIT_SERVER` until the condition preventing the initiate to complete is cleared and `AUDIT_SERVER` initiates correctly. This behavior is deliberate and is designed to prevent the system from running in an incomplete security state.

Conditions that can prevent complete initiation include the following:

- Invalid journal file specification in the audit server database.
- Corrupt Audit Server database (`SYSS$MANAGER:VMS$AUDIT_SERVER.DAT`).
- Corrupt Object Server database (`SYSS$SYSTEM:VMS$OBJECTS.DAT`)
- Failed access to target disk
- Failure to synchronize with other Audit Servers on the cluster
- Required resources being held by other Audit Servers on the cluster

Clearing the condition might require manual intervention. The action required depends on the fault source. Corrective action can include restarting the `AUDIT_SERVER` processes on other cluster nodes or rebooting the affected node in `MINIMUM` state and manually correcting the fault. Corrupt database files can be identified by renaming the files and then restarting the `AUDIT_SERVER`. The server recreates these files and populates them with system default entries.

For more information about booting options, refer to Chapter 4 of the *HP OpenVMS System Management Utilities Reference Manual*.

3.2 C Run-Time Library

Problems corrected in OpenVMS Version 8.3-1H1 include the following:

- The run-time library had faulty code that accessed memory just freed to advance a pointer. In multithreaded code, another thread could reuse that memory before the original thread could advance its pointer. This has been fixed by updating accesses prior to freeing pointers.
- A new processwide exception processing mode—`pure_unix`— has been introduced. In this mode, non-C++ exceptions, also known as OpenVMS conditions, cannot be caught in a C++ catch-all handler. This mode can be requested by calling `cxcl$set_condition(condition_behavior)` with a `pure_unix` argument:

```
cxxl$set_condition(pure_unix);
```

condition_behavior enum declared in <cxx_exception.h> header has been extended to include pure_unix member.

To demonstrate how pure_unix mode works, consider the following program sample. As it is written, the program crashes with ACCVIO. If the call to cxxl\$set_condition() is commented out, the program outputs "caught" and exits:

```
#include <stdio.h>
#include <cxx_exception.h>

void generateACCVIO() { *((int*)0) = 0; }

int main() {
  cxxl$set_condition(pure_unix);
  try {
    generateACCVIO();
  }
  catch(...) {
    puts("caught");
  }
}
```

Note

To use this new functionality you must have a new version of cxx_exception.h, which is included in the CXXL\$ANSI_DEF.TLB file provided with the Version 7.3 compiler (or higher).

- The run-time library sometimes failed to destruct objects of automatic storage duration defined in a function, if such a function exited via an exception that could be caught. This problem has been fixed.
- The run-time library now allows the thread cancel signal (CMA\$_ALERTED) and the thread exit signal (CMA\$_EXIT_THREAD) to be caught in a catch handler with a pointer or a reference to type CXXL\$PTHREAD_CANCEL (or CX6L\$PTHREAD_CANCEL) and CXXL\$PTHREAD_EXIT (or CX6L\$PTHREAD_EXIT), respectively, if catching the signals are enabled. The new types catch these signals exclusively.

Note

To use this new functionality, you must have a new version of cxx_exception.h, which is included in the CXXL\$ANSI_DEF.TLB provided with the V7.3 compiler (or higher).

- The C++ RTL has changed its internal mapping of SIGTRAP from SSS_BREAK to SSS_TBIT, to match a recent CRTL change.
- The C++ RTL used to call std::terminate() when a destructor raised an exception during stack unwinding, even if that destructor did not exit via the exception. This problem has been fixed.
- The C++ RTL used to call std::terminate(), if a foreign exception (such as a non-C++ OpenVMS condition) was raised while a C++ exception was being processed. This behavior has been refined to calling std::terminate() only if the raised OpenVMS condition also leads to unwinding the stack.

OpenVMS Version 8.3–1H1 Release Notes

3.2 C Run-Time Library

- Because OpenVMS conditions can be caught in C++ catch handlers, the C++ RTL converts the conditions to an internal format that matches the representation of C++ exceptions. This conversion would sometimes lead to incorrect information being shown in the traceback. This problem has been fixed.

The following problems are fixed in this version of the C++ Library (Version 7.3 and higher compiler):

- As described in <http://issues.apache.org/jira/browse/STDCXX-397> (the Apache Software Foundation Issues website), the `__introsort_loop()` function in `<algorithm.cc>` header has a bug which, for some input sequences, can adversely affect performance of `std::sort`. For more information, see the Apache tracker for the issue STDCXX-397 at <http://issues.apache.org/jira/browse/STDCXX-397>.

The bug has been fixed. However, for some input sequences, the fix can change the behavior of `std::sort` with regard to the relative order in which elements that have equivalent ordering are placed into the sorted sequence. Though this change in behavior is permissible because, unlike `std::stable_sort`, `std::sort` does not guarantee any particular relative order of elements having equivalent ordering, to avoid breaking applications that rely on existing behaviour of `std::sort`, the fix is conditionalized with `__RW_FIX_APACHE_STDCXX_397` macro. The fix is in effect only when the program is compiled with this macro defined.

- When compiled in standard GNU mode, the library now defines the `_RWSTD_NO_IMPLICIT_INCLUSION` macro, which causes library headers to include their respective template definition files. This is necessary because in standard GNU mode, implicit inclusion is disabled.

Before this change, the program below would link with undefined symbol when compiled in standard GNU mode:

```
#include <vector>

int main() {
    std::vector<int> v;
    v.push_back(0);
}
```

- According to section 27.6.1.3 [lib.istream.unformatted] of the C++ Standard, the following get member functions of the `std::basic_istream` class should call `setstate(failbit)` if no characters have been stored, as is the case for an empty line. While on I64 systems the functions set failbit, on Alpha systems they do not, for example:

```
istream_type& get(char_type *s, streamsize n, char_type delim);
istream_type& get(char_type *s, streamsize n);
```

3.2.1 Ctrl/P Does Not Always Work

On certain Integrity configurations, typing Ctrl/P at the console does not cause OpenVMS to display the IPC (Interrupt Priority C) menu. If you plan to use Ctrl/P, you should test it to ensure that it works, after installing or upgrading to OpenVMS Version 8.3-1H1.

If necessary, you can restore Ctrl/P functionality by performing the following steps:

1. Invoke SDA to analyze the running system:

```
$ ANALYZE/SYSTEM
```

2. Use the CLUE CONFIG to display the adapters on the system:

```
SDA> CLUE CONFIG/ADAPTER
```

3. Locate the "Console Serial Line Driver" adapter (SRA:) in the display:

```
System Adapter Configuration:
```

```
-----
TR Adapter   ADP           Hose Bus   BusArrayEntry   Node GSIN  iVec  SCB
-----
...
5 ACPI_IA64_I FFFFFFFF.8832E0C0  0 00 IA64_BUS
6 PCI         FFFFFFFF.88342A80  9 00 PCI
                                     FFFFFFFF.88342E58  0 0018 00DF 15F0
                                     FFFFFFFF.88342F68  8 0018 00DF 15F0
...

```

```
Port Slot Device Name/HW-Id
-----
```

```
SRA: 0 Console Serial Line Driver
EWA: 1 A6865A (Fast Ethernet)
```

4. Identify the controller that shares the same GSIN (Global System Interrupt Number) as SRA:. In this example, it is EWA:.
5. Exit from SDA and enter the following command (substituting EWA with the correct controller):

```
$ SET DEVICE EWA0/PREFERRED_CPUS='F$GETSYI("PRIMARY_CPUID")'
```

When you complete these steps, Ctrl/P should now function correctly. HP recommends that you edit SYS\$MANAGER:SYLOGICALS.COM to include the SET DEVICE command to ensure correct behavior when the system reboots. Note that Ctrl/P might stop working again if you add or remove I/O adapters. If this happens, redo the steps listed above.

Note also that if XDELTA or the System Code Debugger has been loaded when the system was booted, Ctrl/P is not affected. Entering Ctrl/P will cause the XDELTA prompt to be displayed, for example:

```
Console Brk at 807CF3D2 on CPU 0
807CF3D2!          cmp4.lt      p0, p6 = 3F, r4 (New IPL = 3)
```

HP hopes to provide a fix to this problem shortly with an ECO kit.

3.3 Encryption Issues and Resolutions

The following sections describe corrected and known issues in the Encrypt utility for OpenVMS V8.3-1H1.

3.3.1 PARSE=EXTENDED

Previously, a problem existed with Encrypt version V2.0 where %ENCRYPT-F-KEYLENERR errors were reported if the process parse style was set to extended. This usually occurred when specifying the algorithm with the /KEY_ALGORITHM or the /DATA_ALGORITHM encrypt command qualifiers, which may result in causing a modify access ACCVIO error.

This problem has been corrected in OpenVMS Version 8.3-1H1.

OpenVMS Version 8.3–1H1 Release Notes

3.3 Encryption Issues and Resolutions

3.3.2 Large File Encryption

Previously, a problem existed where Encrypt could not encrypt file larger than 4,194,303 blocks (2,147,483,136 bytes). Encrypt immediately returned to the command prompt with no error message, leading the user to believe that the file was properly encrypted. A CRC error warning would result when attempting to decrypt the corrupted encrypted file.

This problem has been corrected in OpenVMS Version 8.3-1H1.

3.3.3 Input/Output File Performance

A problem has been addressed that corrects Encrypt's I/O performance. This performance was substantially improved (approximately 40x) using multiple buffered reads and multibuffered, multiple writes using \$QIO, eliminating two levels of caching (RMS and XFC) over the single RMS block I/O \$READ and \$WRITE. This results in improved performance, especially with larger files, and improved CPU utilization and less I/O latency between I/Os. This also decreases overall encryption, decryption, and I/O operation completion time.

3.3.4 ENCRYPT/COMPRESS Issue

A problem exists with all versions of Encrypt when using ENCRYPT/COMPRESS with certain file types, notably .BCK and .EXE files. The problem occurs when trying to recover the original (plain-text) data with the DECRYPT command. A key record CRC or another error can occur. The actual encrypted data is incorrect, but no error occurs until decryption.

The problem is potentially serious if the /DELETE qualifier is used when encrypting, because the original file is deleted. The original file can also be lost if the /OUTPUT=file-spec were not used when encrypting and a PURGE command were used.

The /COMPRESS qualifier has been deprecated, disabling its use with Encrypt until HP can resolve this issue. Decryption of compressed files will continue to work normally.

3.4 Documentation Corrections

The following sections describe corrections to various manuals in the OpenVMS documentation set.

3.4.1 HP OpenVMS DCL Dictionary for OpenVMS: N–Z

The *HP OpenVMS DCL Dictionary: N–Z*, Version 8.3 contains two occurrences of a typographical error in the SET SERVER ACME_SERVER, starting on page DCLII-275.

- The first occurrence is in the Description section within the following paragraph:
“In the event of an abnormal process termination, the server will automatically restart and run the SYSSMANAGER:ACMES\$START.COM command procedure. The executive-mode logical name ACMS\$START is used to locate this file.”
The logical name ACMS\$START is incorrectly spelled. The correct logical name is ACME\$START.
- The second occurrence is in the Qualifiers section within the following paragraph:

“The optional keyword, AUTO, causes the server to startup and configure itself using the SYSSMANAGER:ACMES\$START.COM procedure. By default, the server starts with only the OpenVMS ACMS agent configured. The following options are available:”

OpenVMS ACMS agent should read OpenVMS ACME agent.

3.4.1.1 RUN Command Restriction

In the *HP OpenVMS DCL Dictionary*, under the RUN command, the following restriction should be added to the definition for *filespec*:

For I64, the specification for executable image must not include a network (DECnet) or a tape device. Invoking any image over the network or from a tape is not supported on I64.

The following change applies to the *filespec* parameter to the RUN (Image) command:

Specifies an executable image to be executed. The file type defaults to .EXE. The asterisk (*) and the percent sign (%) wildcard characters are not allowed. For I64, the specification for executable images must not include a network (DECnet) or a tape device. Invoking any image over the network or from a tape is not supported on I64.

3.4.2 HP OpenVMS Version 8.3 New Features and Documentation Overview

In the *HP OpenVMS Version 8.3 New Features and Documentation Overview*, Section 3.17, SCD_HARD_OFFFLD should be spelled SCH_HARD_OFFFLD.

Also, because of a technical issue, the SCH_HARD_OFFFLD parameter is no longer used starting in OpenVMS V8.3. See Section 3.1.2 in this document.

3.4.3 HP OpenVMS Version 8.3 Release Notes

References to "Itanium 3" should be changed to "Dual-Core Itanium Processor."

In *HP OpenVMS Version 8.3 Release Notes*, Section 4.19.2, Obsolete System Parameters, the correct system parameters in Open VMS Version 8.3 are the following:

- SMP_CPUS
- SMP_CPUSH
- IO_PREFER_CPUS

The system parameter listed as IO_PREFER_CPU should be IO_PREFER_CPUS.

The SCH_CTLFLAGS parameter, though no longer in Version 8.3, has been renamed SCHED_FLAGS.

In the *HP OpenVMS Version 8.3 Release Notes*, Section 4.19.2, "Obsolete System Parameters," the list of obsolete parameters is incorrect. Table 3–3 lists the parameters.

OpenVMS Version 8.3–1H1 Release Notes

3.4 Documentation Corrections

Table 3–3 Obsolete and Replacement Parameters

Obsolete in OpenVMS V8.3	OpenVMS V8.3 Replacement
SMP_CPUS	SMP_CPU_BITMAP
SMP_CPUSH	-----
IO_PREFER_CPUS	IO_PRCPU_BITMAP

3.4.4 HP OpenVMS System Services

References to "Itanium 3" should be changed to "Dual-Core Itanium Processor."

In the *HP OpenVMS System Services Reference Manual: GETUTC-Z*, under `$$NDOPR`, "Message Buffer Format for `OPCS_RQ_LOGI`," the text, "receive the cancellation request," should read "be affected by the log file initialization."

Under `OPCSL_MS_RQSTID`, references to `OPCSB_MS_TARGET` should be `OPCSZ_MS_TARGET_CLASSES`.

3.4.5 HP TCP/IP Services for OpenVMS Sockets API and System Services Programming Manual

In the *TCP/IP Services for OpenVMS Sockets API and System Services Programming* manual, Table B-2, "end-of-band mark" should be "out-of-band mark."

3.4.6 HP OpenVMS System Management Utilities Reference Manual

In *HP OpenVMS System Management Utilities Reference Manual*, the following Usage Notes should be added to the "Parameter" description of the `CREATE` command:

Usage Notes

- The specified file is expected to be a native image for the platform the install utility is running on.
- If the `CREATE` command is used without a qualifier, the installation utility only creates an entry in the known file list for the specified file. The known file list is not restricted to images; any file can be made known.
- On I64 systems, the specified file can contain an Alpha image. The install utility handles it similar to native images. For example, it creates global sections for the image sections, making the image a shared known image. This is only useful for data-only shareable images. Code sections cannot be executed; main images cannot be run with the `DCL RUN` command. The `LIST` command flags these images with "A*" when displaying the file name of the known image. To avoid installing a wrong image, the system manager should be aware of this feature, especially in a mixed-architecture environment.

The following content should be added to the Parameter description of the `LIST` command:

"On I64 systems, installed Alpha images are flagged with "A*" when displaying the file name of the known image."

3.4.7 HP TCP/IP Services for OpenVMS Management

In the *TCP/IP Services for OpenVMS Management* manual, Section 7.3, "Load Broker Concepts," the following statement is incorrect:

"If a value is specified, the value is used instead of the term (min(235,IJOB LIM) in the following equation."

The correct text is: "If a value is specified, the value is scaled to the effect that the term (min(235,IJOB LIM) * 100) is replaced by (CPU_RATING * 235)."

In the chapter, "Configuring and Managing the NFS Server," the following code example is incorrect:

```
$ @SYS$MANAGER:TCPIP$DEFINE_COMMANDS
$ SYSCONFIG -r nfs_server noproxy_uid=-2 noproxy_gid=-2
```

It should be changed to:

```
$ @SYS$MANAGER:TCPIP$DEFINE_COMMANDS
$ SYSCONFIG -r nfs noproxy_uid=-2 noproxy_gid=-2
```

3.5 GiCAP Functionality Not Currently Available

While running SYSS\$MANAGER:ICAP\$CONFIG.COM, if you respond "Y" to the "Enter (Y)es to configure this system with GiCAP support (N):" prompt, the following message is displayed:

```
HP OpenVMS Industry Standard 64
Global Instant Capacity on Demand (GiCAP) configuration utility
*** GiCAP functionality is not currently available ***
*** GiCAP will be enabled at a later date via an ECO kit ***
```

Also, note that the release notes for iCAP (Instant Capacity) specify GiCAP support for OpenVMS Version 8.3–1H1. This support is not available currently, but will be available in a future update kit. For more information, see the OpenVMS website.

3.6 Iconv Fixes

The following fixes have been made to iconv function/utility:

- Previously, iconv() does not handle NULL parameters on input. This problem has been corrected.
- The iconv converters for DEC Korean can convert to and from Unicode (UCS-2/UCS-4/UCS-8) correctly.

3.7 No-IOLock8 Fibre Channel Port Drivers

Many I/O subsystem components synchronize their operations across CPUs by using the IOLock8 spinlock, which has made acquiring the spinlock a performance bottleneck. As of Version 8.3-1H1, each Fibre Channel port driver (SYSPGQDRIVER, SYSPGADRIVER and SYSPFGEDRIVER) device uses its own port-specific spinlock instead of IOLock8 to synchronize its internal operations. In most configurations, this results in a significant decrease in the amount of time each CPU spends waiting for the IOLock8 spinlock as well as some increase in the Fibre Channel I/O rate.

OpenVMS Version 8.3–1H1 Release Notes

3.7 No-IOLock8 Fibre Channel Port Drivers

Some minor changes are required to any class driver that connects to one of these new port drivers, so customers must determine whether they are running any non-HP class drivers that will not work with them. The simplest way to do this is to examine the output of the SDA command `CLUE SCSI /SUMMARY` and see whether the name of any third-party class driver device appears in the device hierarchy for an FGx0 or PGx0 port device in the "Device" column.

For more information, refer to the notes following this sample SDA session.

```
$ ANALYZE/SYSTEM
OpenVMS system analyzer
SDA> CLUE SCSI /SUMMARY
SCSI Summary Configuration:
-----
SPDT      Port  STDT  SCSI-Id  SCDT  SCSI-Lun  Device  UCB      Type  Rev
-----
81624200  FGB0
          8162CDC0  3
          8162D240  0  GGA22    8162F380  HSV200
          8162F180  1  DGA22801 8162FD40  HSV200  6100
          81632900  2  DGA22802 81632AC0  HSV200  6100
          816354C0  3  DGA22803 81635680  HSV200  6100
          81638080  4  DGA22804 81638240  HSV200  6100
          8162D400  4
          8162DD80  0  GGA22    8163AC40  MRD200
          8163B5C0  1  RJA22801 8163B780  RFD200  6100
          8163C840  2  RJA22802 8163CA00  RFD200  6100
          8163DAC0  3  RJA22803 8163DC80  RFD200  6100
          8163ED40  4  RJA22804 8163EF00  RFD200  6100
```

Note

All of the DGA and GGA devices in this output are accessed through the modified HP class drivers `SYSSDKDRIVER` and `SYSSGKDRIVER` respectively, so they are safe to use with the new port drivers.

Note that even though the physical device of Type `MRD200` is not an HP qualified device, it does not present an `IOLock8` problem because it is accessed through through a `GGAx` unit, indicating that it uses the modified HP Generic class driver `SYSSGKDRIVER`.

The `RJA` devices are not controlled by a modified HP class driver so they will not work with the new port drivers.

3.8 Linker Utility for OpenVMS I64

The following sections describe corrected problems in the Linker Utility for OpenVMS I64.

3.8.1 Linker Writes Incorrect Interimage Debug Fixups into Debug Symbol File

In some situations, the linker creates interimage fixups for the OpenVMS debugger. The inter-image debug fixup is a result of compiler-generated debug relocations, which the linker cannot resolve. That is, the debugger requires these fixups to determine a value from a shareable image at run time. Compilers might differ on how often they require the linker to create interimage fixups for the debugger. The C compiler rarely uses inter-image debugger fixups, although the C++ compiler often uses them. When linking such images with the `/DEBUG`

qualifier, the linker writes the debug information followed by the interimage debug fixups. With the /NODSF qualifier (the default) the information is written correctly into the image file, but with /DSF the information is sometimes written incorrectly into the DSF file.

For example, the DEBUG informationals and the DEBUG error in the following sample occur because the linker has written the DSF file incorrectly:

```
$ RUN/DEBUG MINIREF
%DEBUG-I-INFODWARF, error reading Dwarf info: Section 0a extends outside file
%DEBUG-I-INFODWARF, error reading Dwarf info: Section 0c extends outside file
%DEBUG-I-INFODWARF, error reading Dwarf info: SHT_VMS_FIXUP section 10 size 17eb
e0 not multiple of 18
%DEBUG-I-INFODWARF, error reading Dwarf info: SHT_VMS_FIXUP section 12 size 17ec
30 not multiple of 18
```

OpenVMS I64 Debug64 Version V8.3-003

```
%DEBUG-F-ACCVIO, access violation, reason mask=00, virtual address=000000000014A
000, PC=000000007BD73100, PS=0000001B
%DEBUG-I-INITIAL, Language: C, Module: MINIREF
```

DBG>

The image file is not affected; it can be executed with the RUN command without any problem:

```
$ RUN MINIREF
```

This error is corrected in the OpenVMS Version 8.3-1H1 Linker.

3.8.2 /SELECTIVE_SEARCH Qualifier Might Incorrectly Ignore Transfer Address

If you have an I64 object module containing a transfer address and you include the module in the link operation with the /SELECTIVE_SEARCH qualifier, the linker cannot find its transfer address.

In the following example, the object module (MAIN.OBJ) contains a transfer address but /SELECTIVE_SEARCH qualifier ignores it:

```
$ LINK MAIN/SELECTIVE_SEARCH
%ILINK-W-USRTFR, image USER:[JOE]MAIN.EXE;l has no user transfer address
```

This condition happens only when I64 object modules, meant to provide the program's transfer address, are included in the link operation with the SELECTIVE_SEARCH attribute. The SELECTIVE_SEARCH attribute is given to an object module when you specify the /SELECTIVE_SEARCH qualifier to the object module in the LINK command or in a LIBRARY command.

For example:

```
$ LINK MAIN/SELECTIVE_SEARCH
```

or

```
$ LIBRARY/INSERT LIB.OLB MAIN.OBJ/SELECTIVE_SEARCH
```

This problem manifests itself in one of two ways:

- The linker displays a warning message. This condition occurs if no other object module in the link operation provides a transfer address, weak or otherwise.

OpenVMS Version 8.3–1H1 Release Notes

3.8 Linker Utility for OpenVMS I64

- The linker does not display a message. This condition occurs if other object modules in the link operation provide a transfer address, weak or otherwise. If the linker fails to properly identify the transfer address from the selectively searched object module, it selects one from the other object modules. That transfer address becomes the main entry point for the image. Even though the map file does indicate the incorrect transfer module and transfer address the linker has assigned, the problem might not become evident until you actually run your application.

This error is corrected in the OpenVMS Version 8.3-1H1 Linker.

3.8.3 Maximum Number of Sections

For more than 65280 sections, the ELF format uses an extended numbering scheme, which was not implemented in the linker on OpenVMS I64 Version 8.3. As a result, the number of sections that a single input object module or shareable image can have was limited. Because the linker usually creates the shareable images with sections, this limit also applied when you created shareable images. In that case, ELF sections were used for exporting C++ templates or shared sections. That is, the maximum number of such interfaces in a shareable image had to be fewer than 65280.

In the OpenVMS Version 8.3-1H1 Linker, this restriction is removed. An input file, an object module or a shareable image can have more than 65280 sections.

3.8.4 Incorrect Creation Date of Shareable Images in the Map File

On OpenVMS I64 platforms, shareable images sometimes showed an incorrect creation date in the linker map. The incorrect date shown was usually 3686. This condition occurred when the linker processed the shareable image as an input file and then extracted the date field, which was then shown in the map. The image itself had the correct date that you can see from ANALYZE/IMAGE output.

This error is corrected in the OpenVMS Version 8.3-1H1 Linker.

3.8.5 Demangler Information Look Up Results in Linker Access Violation

In some cases, when the linker tried to look up demangler information in a shareable image that did not contain such information, the linker aborted with an access violation.

This problem is corrected in the OpenVMS V8.3-1H1 Linker.

3.8.6 Incorrect Secondary Messages for the NOGLOSYM Error Message

For the NOGLOSYM error message, the OpenVMS Version 8.3 Linker printed the wrong relocation type and printed some information twice.

This problem is corrected in OpenVMS Version 8.3-1H1.

3.8.7 Incorrect Information for Undefined Symbols

For undefined symbols, a USEUNDEF operation sometimes incorrectly showed information twice for the same reference. The problem occurred when the compiler generated a pair of relocations (LTOFF22X/LDXMOV) for the reference to an undefined symbol.

This problem is corrected in OpenVMS Version 8.3-1H1.

3.8.8 Incorrect UNMAPFIL Error

If a non-ELF input file was encountered, the linker printed an INVLDHDR error message. After an INVLDHDR error, an incorrect UNMAPFIL error message was printed.

This problem is corrected in OpenVMS Version 8.3-1H1.

3.8.9 Max Ident Length Change for Shareable Images in Map

In the linker map, the linker printed up to 14 characters of the ident from object modules and shareable images. (The maximum length of an ident for an object module is not limited; the maximum length of an ident for a shareable images is 15.)

The Version 8.3-1H1 linker now prints up to 15 characters as the maximum of the ident for a shareable image.

3.8.10 Linkage Type Check for Shareable Images

Previously, the linker did not check the type and linkage for symbols from shareable images.

OpenVMS Version 8.3-1H1 Linker now performs this check.

3.8.11 Program Section Attribute ABS Ignored

On I64 systems, you cannot set the ABS attribute to a program section that contains labels to convert its offsets to constants. The linker prints the error message:

```
%ILINK-E-ILLABSPSC, absolute psect <psect-name> has non-zero length (not allowed)
```

The OpenVMS 8.3-1H1 Linker ignores the ABS program section attribute and prints the following informational message:

```
%ILINK-I-PSCATTIGN, psect attribute ABS is not supported for OpenVMS ELF sections, attribute ignored
```

3.8.12 Linker ACCVIOs when FP_MODE Literal Missing From Command Line

In the Version 8.3, the I64 Linker encountered an access violation when the FP_MODE literal was missing on the the command line.

This problem is corrected in OpenVMS Version 8.3-1H1.

3.9 Analyze Utility for OpenVMS

The following sections describe corrected problems in the Analyze Utility for OpenVMS.

3.9.1 Unwind Data Display Problem Corrected

In OpenVMS Version 8.3, the analyze utility (ANALYZE) showed unwind data twice: once as raw data and again as formatted data. If you selected the segment with unwind data (using the segment number or the ALL keyword), ANALYZE displayed the raw data. ANALYZE displayed only the formatted data, however, if you selected the dynamic segment using segment number along with the ALL or the DYNAMIC keyword.

This problem is corrected in OpenVMS Version 8.3-1H1. The formatted unwind data now appears with the segment.

OpenVMS Version 8.3–1H1 Release Notes

3.9 Analyze Utility for OpenVMS

To avoid formatting the same data twice, ANALYZE no longer shows the unwind information with the dynamic segment. To make the selecting of unwind data easier, ANALYZE allows the UNWIND keyword for the /SEGMENT qualifier.

For equivalent output as the former command for unwind information, /SEGMENT=DYNAMIC, use the /SEGMENT=(DYNAMIC,UNWIND) command. Note that if there are multiple segments with unwind data, the /SEGMENT=UNWIND command formats all of them.

3.9.2 Formatted Symbol Vector Correctly Shown in Data Segement

Previously, the symbol vector summary information did not indicate the segment in which the symbol vector resided. The symbol vector was formatted only in the dynamic segment.

This problem is fixed in OpenVMS Version 8.3-1H1. The formatted symbol vector now appears with the data segment in which it is contained. The formatted symbol vector is embedded in data and visible in a dump of the data.

To avoid formatting the same data twice, the symbol vector is no longer shown with the dynamic segment. To make formatting of the symbol vector easy, the SYMBOL_VECTOR keyword is allowed for the /SEGMENT qualifier. When you specify this keyword, the resulting output is only the formatted symbol vector. The surrounding data are not shown. To show and format all of the data, select the segment by number.

To get equivalent output for the former command /SEGMENT=DYNAMIC for symbol vectors, use the /SEGMENT=(DYNAMIC,SYMBOL_VECTOR) qualifier.

The summary information shows the name of the data segment that contains the symbol vector.

3.9.3 Transfer Array Now Formatted in Data Segment

Previously, if you selected the data segment that contained the transfer array (either by segment number or with the ALL keyword), the transfer array was not formatted. Information about the transfer array was shown only in the summary.

This problem is corrected in OpenVMS Version 8.3-1H1. The formatted tranfer array now appears in the data segment.

3.9.4 System Version Array Now Formatted in Dynamic Segment

System version data is in the dynamic segment. Previously, if you selected the dynamic segment (either by segment number, or with the ALL or DYNAMIC keyword), the system version array was not shown. Information about the system version array was only shown in the summary.

This problem is corrected in OpenVMS Version 8.3-1H1. The formatted system version array now appears in the dynamic segment.

3.9.5 Enhancements for the /SEGMENT Qualifier

Enhancements have been made to the /SEGMENT qualifier for dynamic segments. Analyze has been enhanced to accept keywords for the /SEGMENT=DYNAMIC qualifier to provide customized information. The keywords for selectable information are:

- ALL—(Default) Formats all parts of the dynamic segment
- TAGS—Formats the tag array
- IMAGE_STRINGS—Formats strings of the specified image

- RELOCATIONS—Formats the image relocations
- FIXUPS—Formats the image fixups
- SYSTEM_VERSION_ARRAY—Formats the system version array

The default, /SEGMENT=ALL, formats all of the image information.

Note that formatting using the TAGS keyword includes the names of the needed images, so you do not have to add IMAGE_STRINGS to print the names.

3.9.6 Support for Section Escaping Added

On OpenVMS V8.3, the Analyze utility did not complete when analyzing an object module with more than 65,280 sections. Instead, it looped when attempting to print the section header table.

This problem has been fixed in OpenVMS V8.3-1H1.

3.10 INSTALL Utility for OpenVMS (Installing Resident Images in S2 Space)

The INSTALL utility now supports installing code segments of resident images into 64-bit S2 address space. Not all code can run in a full 64-bit address space (P2 or S2). For example, the code must be prepared for 64-bit PCs when handling exceptions. Also, some compilers require the /POINTER_SIZE=64 command qualifier, when generating code, suitable for a 64-bit address space.

To avoid mapping unprepared code in S2 space, the INSTALL utility by default will continue to map the code segments in S0/S1 space. The INSTALL utility will map code segments of resident images to S2 if two conditions are met:

- The developer explicitly confirmed that the code is 64-bit ready by using the /SEGMENT_ATTRIBUTE=CODE=P2 qualifier when linking the image.
- There is enough pre-allocated space in the resident code region in the S2 space to map the code segments. The size of the region is determined by the system parameter GH_RES_CODE_S2 (number of pages). The default value is set to 0. That means that by default even 64-bit ready resident images have their code mapped in S0/S1 space.

3.11 Librarian Utility

The Librarian Utility now has the following enhancement.

3.11.1 Support for Section Escaping

Previously, object modules with more than 65,280 sections could not be added to an object library. The Librarian aborted with an error message:

```
$ LIBRARIAN/CREATE MYLIB 64K_SECTIONS
%LIBRAR-E-REPLACERR, error replacing USER$DISK:[JOE]64K_SECTIONS.OBJ;1
in USER$DISK:[JOE]MYLIB.OLB;1
$
```

This condition is corrected in OpenVMS Version 8.3-1H1. Support for section escaping has been added for OpenVMS Version 8.3-1H1.

3.12 InfoServer Utility and FDDI

3.12 InfoServer Utility and FDDI

Using the InfoServer utility on OpenVMS to boot a client over an FDDI network adapter is not supported.

3.13 New Qualifier for DCL Command SET PASSWORD

The DCL command SET PASSWORD now accepts the /PROMPT qualifier with two permitted values: /PROMPT=FIXED and /PROMPT=VARIABLE. If you use the SET PASSWORD command in a DCL command procedure, do not specify the /PROMPT=VARIABLE qualifier. If you do, it works as expected, but any failing status is only displayed and not returned to DCL.

3.14 INITIALIZE/ERASE=INIT Before Using Media

HP recommends that you issue the DCL command INITIALIZE/ERASE=INIT on storage media prior to using them for the first time. This eliminates any stale data that was left from previous use by another operating system or diagnostics.

An indication of such stale data is three question marks (???) in the console command output, as shown in the following example:

```
Shell> ls fs1:\
Directory of: fs1:\
 00/00/07 19:16p    1,788,984,016 ???
 00/00/80 12:00a         0 ???
      2 File(s)  1,788,984,016 bytes
      0 Dir(s)
```

The problem will be corrected in a future release.

3.15 XML-C Product Zip File

The XML-C product for OpenVMS for Integrity servers is delivered as a ZIP file that contains a self-extracting executable file. The XML-C installation documentation describes how to install the product by using this executable file. To obtain the executable file, extract it from the ZIP file.

3.16 Using the OpenVMS e-Business and Integration Infrastructure Package

The OpenVMS e-Business and Integration Infrastructure Package for OpenVMS is contained on two CDs that are formatted so that they appear as a Files-11 file structure to an OpenVMS system and an ISO 9660 file structure to a Windows, Linux, or UNIX system.

Installation

The component installation kits and documentation are split across the two CDs. Component installation can be done only on an OpenVMS Alpha system from the specific CD designated in the top-level `index.html` file.

Documentation

For OpenVMS systems, partial component documentation is viewable based on which CD is mounted for use. Component documentation is available only for the components present on the specific CD.

3.16 Using the OpenVMS e-Business and Integration Infrastructure Package

For Windows, Linux, or UNIX systems, complete component documentation is viewable on both CDs.

3.17 Performance Data Collector for OpenVMS (TDC)

TDC_RT Version 2.2-107 is included in the OpenVMS Version 8.3–1H1 installation. An update to TDC Version 2.2-108 is now available from the TDC Web site at:

<http://www.hp.com/products/openvms/tdc/>

TDC Version 2.2-108 corrects several issues discovered in TDC_RT Version 2.2-107. It also enables collection of internet metrics in TCPware and MultiNet environments, adds additional metrics to several data records, and provides new programming features and sample code in the TDC Software Developers Kit.

3.18 C++ Compiler

C++ Version 7.2 for OpenVMS for Integrity servers predefines the macro `__INITIAL_POINTER_SIZE` to 0, unlike C++ Version 7.1 compiler, which leaves it undefined. This is an intentional change that makes C++ Version 7.2 consistent with the C compiler and provides support for `pointer_size` pragmas, while C++ Version 7.1 does not.

Note that this change can cause diagnostics to appear in code that compiled cleanly with certain declarations selected by system header files that declare pointer types. This effect is most likely to appear in applications that use starlet headers and that compile with `__NEW_STARLET` defined.

If you cannot modify the application source code to conform to the new declarations, add the command-line qualifier `/UNDEF=__INITIAL_POINTER_SIZE` to the CXX command line to prevent the C++ Version 7.2 compiler from predefining this macro and thus causing the system headers to provide the same declarations as with Version 7.1 of the compiler.

3.19 Building DCE IDL C++ Applications

Building DCE IDL C++ applications on CXX Version 7.2 and higher results in an undefined symbol linker warning. This is a known issue. To overcome this warning, contact HP Support Services to request any necessary patches.

3.20 CDSA Errors During Installation

You might see the following errors when upgrading to OpenVMS Version 8.3-1H1. You can safely ignore these errors because CDSA and Secure Delivery will still function properly and PCSI kits will install normally and securely.

```
CDSA-I-Init, CDSA has previously been initialized on this system.  
CDSA-I-Init, Re-initializing CDSA.
```

OpenVMS Version 8.3–1H1 Release Notes

3.20 CDSA Errors During Installation

```
CDSA-I-Init, Initializing CDSA
MDS installed successfully.
Module installed successfully.
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module: MDS Error (common-DL): 300A
  (Code #300A)!
Module installed successfully.
Module: MDS Error (common-DL): 300A
  (Code #300A)!
CDSA-I-Init, CDSA Initialization complete CDSA-I-Init, Initializing Secure Delivery.
Install completed successfully.
Module: MDS Error (Clean): 300A
  (Code #300A)!
Module installed successfully.
CDSA-I-Init, Secure Delivery Initialization complete
```

3.21 Update to FREEWARE README.TXT File

An update to the [FREEWARE]FREEWARE_README.TXT, included on each OpenVMS Freeware CD, is available for download from the OpenVMS Freeware Web site at:

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

This updated file includes the correction to the displayed version number from V7.0 to the intended and expected V8.0, as well as additional updates and corrections.

As is traditional with the OpenVMS Freeware, all updates to existing files and new packages are available at the OpenVMS Freeware website.

3.22 Process/Application Hangs

The following restriction applies to the the LIBRTL documentation for the lib\$find_image_symbol run-time library routine:

If your application might dynamically activate shareable images that use pthreads (or the older CMA thread interface), the main image must be linked with the pthread\$rtl image.

3.23 Provisioning

The following release notes pertain to provisioning on OpenVMS Version 8.3-1H1.

3.23.1 BL860c System Must Be Powered ON

When provisioning a BL860c server blade, the system must be powered on. If the system is powered off the upgrade/install will fail, even though you can access the iLO.

3.23.2 InfoServer Name Length

The InfoServer name must be less than 12 characters long for provisioning to work. This is a temporary restriction.

3.23.3 Provisioning Multiple Servers

HP SIM provisioning (InfoServer or vMedia) can provision multiple servers simultaneously, but when vMedia is used independently (without HP SIM), only one server at a time can be provisioned.

3.23.4 Provisioning from HP SIM Central Management Server

OpenVMS can be provisioned from the HP SIM Central Management Server running Microsoft Windows 2003 server or SUSE Linux.

3.23.5 OpenVMS InfoServer and the Server Blade on the Same Local Area Network

The OpenVMS InfoServer and the server blade must be on the same Local Area Network to provision the server blade.

3.23.6 EFI Firmware

The EFI firmware for the BladeSystem must be at version 5.0 or later.

3.23.7 Management Processor

The Management Processor must be running the Advanced iLO2 firmware.

3.24 SYS\$LDDRIVER Restriction

SYS\$LDDRIVER.EXE is a freeware pseudo device driver that allows OpenVMS operating system to create virtual disks. For OpenVMS V7.3-1 and succeeding versions, this driver was placed in SYS\$COMMON:[SYS\$LDR] to support the creation of the source virtual disk for mastering a CD or DVD using CDRECORD or COPY/RECORDABLE_MEDIA. This is the only supported use of this freeware driver. All other uses of this driver continue to be subject to the following documented freeware usage restrictions:

OpenVMS Freeware is provided as is without a warranty. HP imposes no restrictions on its distribution or redistribution. HP does not support services for this software, fix the software, or guarantee that it works correctly.

3.25 RMS \$PARSE Validation of Directory Files

Starting with OpenVMS Version 8.3, the \$PARSE service further validates all directories named in a directory specification to ensure that the directory characteristic is set. In previous OpenVMS versions, attempting to use a file with a .DIR extension that was not a directory resulted in a SSS_BADDIRECTORY error from the \$OPEN service, but not necessarily from the \$PARSE service. As of Version 8.3, the error is consistently returned by the \$PARSE service as long as it is not a syntax-only \$PARSE.

3.26 SCACP Error Counter Can Report Retransmit Errors

If the PEA0: device on the system shows a number of errors, these errors might be retransmit errors and not actual errors. To verify actual errors, use the SCACP utility to confirm whether there are a number of retransmits on the PEA0 channels and use the LANCP utility to identify whether any actual devices errors exist on the LAN devices that the PEdriver uses. If there are retransmits and no devices errors, then the PEA0: device errors are likely retransmits and not actual errors.

3.27 Serial Port Enumeration

On OpenVMS systems, enumeration is the process by which devices are assigned a letter and number following the OpenVMS generic device-type nomenclature. In the case of serial ports, enumeration is expressed as TTA0, TTBO, and so on, for generic serial port devices, and as OPA0 for a serial port device that has been selected as the system's primary console at the EFI Boot Manager or EFI Shell> prompt.

OpenVMS Version 8.2 consistently enumerated system serial ports according to the rules and precedents established by OpenVMS for Alpha systems. With OpenVMS Version 8.3, those rules were violated and users experienced inconsistent port naming, particularly on systems migrating from Version 8.2 to Version 8.3.

OpenVMS Version 8.3-1H1 returns to the consistent serial-port naming conventions established for HP Integrity in OpenVMS Version 8.2, with the goal of not changing serial port names more than necessary, and for consistency with the policy on OpenVMS Alpha systems. The names of the serial ports can change, because at least one serial port can serve more than one function.

The serial port selected as primary console is always OPA0. If the graphics console has been selected as primary, the keyboard and graphics head constitute OPA0, and the serial ports will be named TTA0, TTBO, and so on.

Unless the serial port of the Integrated Lights Out (iLO) Management Processor (MP) is selected as the primary console, it is not connected as a serial port and is not exposed by the Operating System. It is not suitable for general-purpose use because it cannot support the data rates a general-purpose serial port needs to support. This is an optional component in most systems. Check the options list shipped with your system and your system's documentation at the HP Documentation website:

<http://docs.hp.com>

There are two possible serial ports that can be selected as the primary console, the iLO and the Baseboard Management Console (BMC). Whichever is selected as primary console will be expressed as OPA0 by OpenVMS; the other will be either TTA0 or TTBO if the system has an additional serial port. The following list describes these abbreviations and their definitions.

Abbreviation	Definition
MP	Serial port of the iLO MP. This component is optional on some systems.

OpenVMS Version 8.3–1H1 Release Notes

3.27 Serial Port Enumeration

Abbreviation	Definition
BMC	Serial port of the BMC. This component is not available on all systems.
AP	Auxiliary Port. This auxiliary 16550-compatible serial port is not available on all systems.
VGA	Graphics Console. This is an optional component of the iLO MP. If your system was not shipped with the VGA option, you can install a graphics option in one of the PCI slots to obtain this functionality.
NA	Not Available.
NC	Not Configured as a Serial Port by OpenVMS.
NS	Not Supported.

The following table displays the sources for backpanel drawings covered in this section.

Platform	Backpanel Drawing
rx1600, rx1620	http://docs.hp.com/en/AB430-96004/ch03s03.html#i1021437
rx2600, rx2620	http://docs.hp.com/en/AD117-9003A/ch02s02.html
rx4640	http://docs.hp.com/en/A9950-96009/A9950-96009.pdf
rx3600, rx6600	http://docs.hp.com/en/AD217-9001A/ch02s03.html
rx2660	http://docs.hp.com/en/AD217-9001A/ch02s02.html
rx8620	http://docs.hp.com/en/A7026-96033/ch04s05.html
BL860c	http://docs.hp.com/en/AD217-9001A/ch02s01.html

The following table provides serial-port naming for the HP Integrity platforms listed. The device selected as primary console is always named OPA0.

OpenVMS Serial Port Name					
Platform	Primary Console Port	MP	BMC	AP	VGA
rx1600 rx1620	MP (optional) BMC VGA (optional)	OPA0 NC NC	TTA0 OPA0 TTA0	NA	OPA0
rx2600 rx2620	MP (optional) BMC VGA (optional)	OPA0 NC NC	TTB0 OPA0 TTB0	TTA0 TTA0 TTA0	OPA0
rx4640	MP VGA (optional)	OPA0 NC	NA	NA	OPA0
rx3600 rx6600	MP (optional) BMC VGA (optional)	OPA0 NC NC	TTA0 OPA0 TTA0	OPA0	
rx2660	MP VGA (optional)	OPA0 NC	NA	TTA0 TTA0	OPA0
rx8620	MP VGA	OPA0 NC	NA	NA	OPA0
BL860c	MP VGA	OPA0 NC	NA	NA	OPA0

3.28 Old Firmware Cannot Translate Messages Written to the System Event Log

Upon installation, OpenVMS Version 8.3-1H1 begins writing new messages to the system event log. To see the system event log, select it (on most systems) from the main MP menu (SL: Show Event Logs).

Older firmware translates the messages as "IPMI Type-E0 Event," instead of the correct OS_OPENVMS_BUGCHECK and OS_OPENVMS_SHUTDOWN.

The following shows the OS_OPENVMS_BUGCHECK message (alert level *5 - critical) on a system running older firmware:

```
291      0 *5 0xB4801C9700E01B50 000000000019000C IPMI Type-E0 Event
                                     30 Jul 2007 14:03:41
```

The following is an example of OS_OPENVMS_SHUTDOWN message (alert level 2 - informational) on a system running older firmware:

```
296      0 2 0x54801C9900E01BD0 00000000001A000C IPMI Type-E0 Event
                                     30 Jul 2007 14:22:06
```

The new firmware uses the phrase "OS_OPENVMS_BUGCHECK" or "OS_OPENVMS_SHUTDOWN" in place of "IPMI Type-E0 Event".

A third message, OS_BOOT_COMPLETE, has a different alert level on a system running new firmware. It has been changed by OpenVMS to informational, or level 2:

```
301  OS  0 2 0x548016E100E01B80 0000000000000001 OS_BOOT_COMPLETE
                                     23 Aug 2007 14:25:44
```

New firmware displays the following message when "T - View Mode Configuration Text" is selected:

```
MP:SL (+,-,CR,D, F, L, J, H, K, T, A, U, ? for Help, Q or Ctrl-B to Quit) >t .
```

```
.
```

```
Log Entry 301: 23 Aug 2007 14:25:44
```

```
Alert Level 2: Informational
```

```
Keyword: OS_BOOT_COMPLETE
```

```
OS Boot Complete
```

```
Logged by: O/S Kernel (Generic) 0
```

```
Data: Major change in system state - Boot Complete 0x548016E100E01B80 0000000000000001
```

3.29 AST Delivery Clarification in Programs using POSIX Threads

It is possible to utilize ASTs in threaded programs. Section B.12.5 in the Guide to the POSIX Threads Library describes some general usage notes and cautions. However, that section does not make clear how AST delivery behaves in programs with upcalls disabled (which is the default configuration).

In a program with upcalls disabled, user-mode ASTs will interrupt whatever thread happens to be executing at the moment that the AST is delivered. Therefore the AST service routine cannot make any assumptions about the context in which it executes (with respect to thread ID, stack space available, etc.).

OpenVMS Version 8.3–1H1 Release Notes

3.29 AST Delivery Clarification in Programs using POSIX Threads

Also, note that much of the material in Section B.12.5 of the Guide describes a possible future version of OpenVMS. The description of generalized "per-thread" or thread-targeted ASTs represents possible future enhancements to the operating system. In all OpenVMS releases to date, however, user-mode ASTs are treated as if they are directed to the process as a whole.

3.30 TZ Function in the CRTL

The TZ logical name or DCL symbol is used by the C Run-Time Library to define the time zone to be used in certain C program time-related functions. (For a full description of TZ, its use, and specific functions see the C Run-Time Library documentation.)

The TZ logical name or DCL symbol has been used by the C Run-Time Library since Version 7.3 of OpenVMS. However, with Version 8.3, there has been a change.

Prior to Version 8.3, defining TZ to something other than a valid time zone caused the time zone to default to localtime (that is, the current time zone of your system). With OpenVMS Version 8.3, defining TZ to an invalid timezone causes the C Run-Time Library's Time functions to resort to Coordinated Universal Time (UTC) time.

Be aware that if you define the logical name or DCL symbol TZ to a non-standard definition, doing so might cause undesirable side-effects in some C programs.

3.31 HP TCP/IP Services for OpenVMS

HP strongly recommends that you apply the current Engineering Change Order (ECO) for TCP/IP Services for OpenVMS Version 5.6 after upgrading to OpenVMS Version 8.3–1H1 to ensure that the latest software is installed.

The following information applies to the current release of HP TCP/IP Services for OpenVMS V5.6.

3.31.1 NI_* Flag Values for the getnameinfo() function Improperly Changed

The NI_* flag values for the getnameinfo() function were improperly changed with the V5.6 release. This would cause any applications using the NI_* flag values that were built against pre-V5.6 TCP/IP versions to not run as expected on TCP/IP V5.6. This problem has been corrected and the flag values have been returned to their pre-V5.6 definitions. Note that any applications using the NI_* flag values that were built against V5.6 will no longer execute properly on V5.6 ECO1 or later. These applications should be rebuilt.

Note that when building applications against TCP/IP V5.6 ECO1, the NI_* constants defined in netdb.h in sys\$common:[syslib]decc\$rtldef.tlb may not correspond to the values in the TCP/IP product. The NI_* constants in netdb.h in decc\$rtldef.tlb will be updated with the next C RTL release of that library, but due to the separate release of products, they may be out of synch for a period of time. For V5.6 ECO1, the NI_* constants are defined as follows:

```
#define NI_DGRAM          0x00000001
#define NI_NAMEREQD      0x00000002
#define NI_NOFQDN        0x00000004
#define NI_NUMERICHOST   0x00000008
#define NI_NUMERICSERV   0x00000010
#define NI_NUMERICSCOPE 0x00000020
#define NI_WITHSCOPEID   0x00000040
```

OpenVMS Version 8.3–1H1 Release Notes

3.31 HP TCP/IP Services for OpenVMS

If your application uses these constants, you may need to manually update the `netdb.h` file. Alternatively, you may update the `netdb.h` that resides in the `tcpip$examples` directory and then `#include "tcpip$examples:netdb.h"`.

Additional changes that have been made to `netdb.h` for V5.6; therefore, you should not use an older version of the `netdb.h` file in place of the V5.6 file.

3.32 Virtual Connect

The following sections pertain to Virtual Connect.

3.32.1 Failover and `RECNXINTERVAL`

`RECNXINTERVAL` may have to be increased above the default of 20 to allow time for Virtual Connect Manager failovers. This is especially true in larger clusters.

3.32.2 Fibre SCSI

If you are installing OpenVMS Version 3.1-1H1 on an HP Integrity BL860c Server Blade, an updated Fibre SCSI kit is required for Virtual Connect. Refer to the OpenVMS web site for updated kits:

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

3.33 Sysgen Parameter `CPU_POWER_MGMT` Default Has Changed

The default value for the sysgen parameter `CPU_POWER_MGMT` has changed from 1 to 0 (that is, from on to off). This change improves the latency of interrupts and of CPUs exiting from IDLE. This value may change again in future versions of OpenVMS.

3.34 Booting A Satellite System with Reserved Memory

To use the `SYSMAN` reserved memory feature on an Integrity satellite system the file `SYSS$SYSTEM:VMS$RESERVED_MEMORY.DATA` must allow world `READ+EXECUTE` access. Failure to set this access protection results in the warning when booting the satellite:

```
%VMS_LOADER-W-Warning: Unable to load file SYSS$SYSTEM:VMS$RESERVED_MEMORY.DATA
```

After running `SYSMAN` to add memory reservations to a satellite, execute `SYSS$MANAGER:CLUSTER_CONFIG_LAN.COM` to set the correct protection on the `VMS$RESERVED_MEMORY.DATA` file. To set the protection, from the cluster configuration procedure "Main Menu" select:

3. CHANGE a cluster member's characteristics.

From the "CHANGE Menu" select the following:

13. Reset an IA64 satellite node's boot environment file protections.

What is the satellite name (leave blank to use a specific device and root)?

Enter the satellite name or satellite boot device and root for the system where you added the memory reservation. `SYSMAN` will be fixed in a later release to eliminate this condition.

MSA Utility

This chapter contains information about the MSA utility.

4.1 MSA Utility Usage Summary

The MSA utility is an OpenVMS system management tool for configuring and managing the following controllers:

- HP StorageWorks Smart Array Family of Storage Host Bus Adapters (5300 series, 6400 series, and P400)
- HP StorageWorks Modular Smart Array Family of Storage Controllers (MSA1000 and MSA1500)

These controllers connect to HP StorageWorks Modular Smart Array storage array systems, and to Integrity Server internal SCSI or SAS drives.

4.1.1 Required Privileges

You must have the following privileges to run the MSA utility: NETMBX, TMPMBX, SYSPRV, DIAGNOSE, PHY_IO.

To start the MSA utility, enter the following command at the DCL command prompt (\$):

```
$ RUN SYS$SYSTEM:MSA$UTIL
```

The MSA utility returns the following prompt:

```
MSA>
```

At the MSA prompt, you can enter any MSA utility command described in the following sections.

4.1.2 Restrictions

The MSA Utility has the following restrictions:

- The MSA_UTIL command SHOW CONNECTIONS is not applicable to or supported on the HP StorageWorks Smart Array Controller.
- EXTEND, MIGRATE, and EXPAND commands on Smart Array systems will fail if the controller cache is bad or the cache battery is less than 75%. These commands work on controllers that are designed to support volume expansion by using disk drives for backing up expand data stripe. This is in the absence of battery-backed memory on the controller.
- RAID units that have spare disks (RAID 1, RAID 5, and so on) cannot be migrated to RAID 0 or JBOD units.
- Multiple-capacity disks can be used by specifying the /SIZE qualifier. The maximum value of size specified with the /SIZE qualifier should be the size available on the disk with the lowest capacity.

MSA Utility

4.2 MSA Utility Commands

4.2 MSA Utility Commands

The following sections describe MSA utility commands and provide examples of their use. Note that some qualifiers are shown as abbreviated in the examples.

ACCEPT UNIT

Changes the state of the unit back to VOLUME_OK when all drives of a previously failed unit are in working order.

This command accepts media exchange on a unit marked as failed.

Note

The ACCEPT UNITS command resets the status of all failed units to VOLUME_OK.

Format

ACCEPT UNIT <#>

where # represents the unit number.

Example

```
MSA> ACCEPT UNIT 2
```

ADD UNIT

Creates units (logical storage units that comprise one or more hard drives).

Format

ADD UNIT <unit_n> / <qualifiers>

Parameters

unit_n

The unit number can range from 0-31.

Note

Any two or a combination of all the following qualifiers cannot be used in conjunction with each other.

Qualifiers

/ADG

Specifies the RAID type for the unit as Advanced Data Guard (ADG).

/DISK

Specifies the disk numbers to be used to form the unit. Enclose multiple disks in parentheses. Use the following format:

/DISK=(disk-numbers[, ...])

/IDENTIFIER

User-defined identifier for the unit. This identifier is used by OpenVMS to name the device. The value of the identifier *n* is between 1 and 9999. Use the following format:

/IDENTIFIER= n

Note

The /IDENTIFIER qualifier is not required for Smart Array controllers.

/JBOD

Specifies the RAID type as JBOD. Synonymous with RAID 0.

/PARTITION

Specifies the partition number to be used for a given unit. The first unit that is created on a disk/disk group is automatically assigned the partition number 0. Units that are subsequently created on this disk/disk group must be created with sequential partition numbers. Use the following format:

/PARTITION=(partition_number)

/RAID_LEVEL

Specifies the RAID type of the unit. Supported values for this qualifier are 0 (data striping), 1 (data mirroring), and 5 (data striping with striped parity).

Use the following format:

```
/RAID=[(0 | 1 | 5)]
```

/SIZE

Specifies the size of the unit. If the size qualifier is not specified, the size defaults to the maximum capacity of the disks depending on the RAID level, as shown in the following example:

```
/SIZE=#(GB | MB | KB | %)
```

/SPARE

Specifies the disk numbers to be used as spare disks. Enclose multiple disks in parentheses. Assigning a spare disk to an unit in a drive group assigns the spare disk to all the configured units in the drive group.

If a unit is created on a disk group to which a spare disk is assigned, then the spare disk is configured to the new unit (if it is not a RAID 0 unit). One spare disk can be assigned to multiple drive groups.

Ensure that the size of the spare disk is equal at least to the size of the smallest drive in the drive group. Use the following format:

```
/SPARE=(disk_number[,...])
```

/STRIPE_SIZE

Specifies the stripe size for a given RAID volume. Stripe size must be one of the following values: 8, 16, 32, 64, 128, or 256. Raid 5 and ADG are limited to a maximum 64 KB stripes. RAID 0 and 1 default to 128 KB stripes, and RAID 5 and ADG default to 16 KB stripes. Use the following format:

```
/STRIPE_SIZE=(stripe_size)
```

/VERBOSE

Provides logging that can be interpreted by engineering.

Examples

1. MSA> ADD UNIT 2

This command creates unit 2.

2. MSA> ADD UNIT 3/ID=1003/DISK=103/JBOD/PARTITION=0/SIZE=8GB
MSA> ADD UNIT 4/ID=1004/DISK=103/JBOD/PARTITION=1/SIZE=10GB
MSA> ADD UNIT 5/ID=1005/DISK=103/JBOD/PARTITION=2/SIZE=8GB

These commands create three units on disk 103. Notice that the partition numbers are provided in sequential order on the same disk/disk group.

3. MSA> ADD UNIT 3/ID=1003/DISK=103/JBOD/PARTITION=0/SIZE=8GB
MSA> ADD UNIT 4/ID=1004/DISK=103/JBOD/PARTITION=2/SIZE=10GB

In this command sequence, assume that there is no unit with partition number 1 on this disk, the creation of unit 4 will fail because the unit is assigned the partition number 2 and the assignment is not in sequential order.

MSA Utility DELETE UNIT

DELETE UNIT

Delete units from the drive. Regardless, whether the disk is mounted on OpenVMS, you are prompted for confirmation before proceeding to delete the selected unit.

For example, if the disk is mounted on an OpenVMS node, in addition to the prompting for confirmation, you are informed that the disk is mounted.

Note

After a unit is deleted, its unit number remains unused until manually assigned to a new unit.

Unit numbers are not automatically reassigned when a unit is deleted.

If more than one unit has been created on a disk/disk group, only the last created unit can be deleted.

Be sure to maintain a record of the unit numbers and the order in which they are created on a disk/disk group.

Format

```
DELETE UNIT <unit_n> / <qualifiers>
```

Parameters

unit_n

The unit number can range from 0-31.

Qualifiers

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> DELETE UNIT 4
```

In this example, unit 4 is the unit to be deleted. This is the same number that was given to the unit when it was created using the ADD UNIT command.

EXIT

Exits the MSA\$UTIL program.

MSA Utility FLASH FIRMWARE

FLASH FIRMWARE

Updates the firmware of a given controller. Specify the firmware file name to be used.

Format

FLASH FIRMWARE <firmware-file-name> / <qualifiers>

Parameters

firmware-file-name

Qualifiers

/VERBOSE

Provides logging that can be interpreted by engineering.

HELP

Describes all currently supported MSA\$UTIL commands, their parameters, and their qualifiers. Use the following format, where *verb* is the specific description:

Format

HELP <verb>

Parameters

verb

Example

```
MSA> Help ADD
ADD
UNIT
    ADD UNIT is used to create UNITS (logical storage units comprising
    one or more hard drives).
Format:
ADD UNIT <unit_n> <qualifiers>
Parameters  Qualifiers  Examples
```

This command describes the ADD command and its parameters.

LOCATE

Causes the LEDs of the requested drives to blink. These LEDs are visible from the front of the attached storage enclosures.

Note

If a time limit is not included with the LOCATE command, the LEDs blink for 30 seconds.

Format

LOCATE/TIME=xxx <parameters> / <qualifiers>

Parameters

ALL

Causes all drives connected to the MSA storage subsystem to flash.

BUS *bus-number*

Causes the disks connected to the specified bus to flash.

CANCEL

Cancels the current locate operation.

DISK *disk-number*

Causes the specified drive to flash. Disks are identified by their bay number.

UNIT *unit-number*

Causes the disks configured on the specified unit to flash.

Qualifiers

/TIME=time

Specifies the number of seconds the disk's LED should flash. This is an optional qualifier whose default is 30 seconds.

/VERBOSE

Provides logging that can be interpreted by engineering.

Examples

1. MSA> LOCATE ALL

This command locates all drives attached to the storage enclosure.

2. MSA> LOCATE BUS 1

This command locates all drives SCSI bus 1.

3. MSA> LOCATE UNIT 1

This command locates all drives in unit 1.

4. MSA> LOCATE DISK 102

This command locates drive 102.

MSA Utility READ FIRMWARE

READ FIRMWARE

Reads the firmware image on the controller.

Note

Supported only for Smart Array 5300 and Smart Array 6400 series controllers.

Format

READ FIRMWARE <firmware file-name> / <qualifiers>

Parameters

firmware file name

Specifies the file name of the firmware image.

Qualifiers

/VERBOSE

Provides logging that can be interpreted by engineering.

RESET THIS_CONTROLLER

Issues a controller reset to the specified controller.

Format

RESET THIS_CONTROLLER <parameters> / <qualifiers>

Parameters

CONTROLLER_TYPE

Issues a controller reset to the specified controller.

Note

Applicable only for MSA1000/1500 controllers.

Qualifiers

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> RESET THIS_CONTROLLER
```

MSA Utility RESET OTHER_CONTROLLER

RESET OTHER_CONTROLLER

Issues a reset to the controller. The state of this controller can be either ACTIVE or STANDBY.

Format

RESET OTHER_CONTROLLER / <qualifiers>

Parameters

CONTROLLER_TYPE

Issues a controller reset to the specified controller.

Note

Applicable only for MSA1000/1500 controllers.

Qualifiers

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> RESET OTHER_CONTROLLER
```

SCAN ALL

Sends a scan message to the Smart Array controller instructing it to scan SCSI buses and discover new or replaced disks. After the scan is complete, the rebuild operation for the logical volumes is initiated for all the units configured in the disk group.

This command is applicable only for the internal disk enclosure connected to the Smart Array on I64 platforms.

Note

For additional information on rebuild of volumes, see the START RECOVER command.

Format

SCAN ALL

Parameters

None

Qualifiers

None

Example

```
MSA> SCAN ALL
```

SET CONTROLLER

Selects the controller device-name as the default controller. The device name has the format *ddcu:*, where:

dd is the device code.

c is the controller designation (A through Z).

u is the unit number (0 through 9999).

This command is required before entering all SAS utility commands except SHOW CONTROLLER and SHOW VERSION.

Note

SET CONTROLLER command is synonymous with the SET ADAPTER command.

Format

SET CONTROLLER [ddcu:] / <qualifiers>

Parameters

None

Qualifiers

/DEFAULT

Specifies the default controller.

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> SET CONTROLLER $1$GGA105:
```

This command sets controller \$1\$GGA105 as the default controller.

SET GLOBALS

Set the expand priority, the read/write ratio, and the system name.

Format

SET GLOBALS / <qualifiers>

Parameters

None

Qualifiers

/EXPAND_PRIORITY=(low | medium | high)

Sets the default controller.

Specifies the expand priority. Use when expanding an array to set the priority of array expansions in relation to input/output (I/O) operations. Use in the following format:

/EXPAND_PRIORITY=(LOW | MEDIUM | HIGH)

/READ_CACHE=value

Specifies the percentage of cache to be used with the READ command.

This value must be between 0 and 100. Use the following format:

/READ_CACHE=(AUTOMATIC | percent)

Note

Combined, the values of the /READ_CACHE and the /WRITE_CACHE qualifiers must equal 100.

/REBUILD_PRIORITY=(LOW | MEDIUM | HIGH)

Specifies the rebuild priority of the RAID volume.

Use when rebuilding an array to set the priority of an array rebuild in relation to input/output (I/O) operations. Low expansion or rebuild applies only when the array controller is not busy handling normal I/O requests. This setting has minimal effect on normal I/O operations. However, there is an increased risk that data will be lost if another physical drive fails while the rebuild is in progress. Use the following format:

/REBUILD_PRIORITY=(LOW | MEDIUM | HIGH)

/WRITE_CACHE=value

Specifies the percentage of cache to be used with the WRITE command.

MSA Utility SET GLOBALS

The value must be between 0 and 100.

Note

Combined, the values of the /READ_CACHE and the /WRITE_CACHE qualifiers must equal 100.

/SYSTEM_NAME=*name*

Specifies the name to be assigned to the controller.

The name represents any-user defined phrase, up to 20 alphanumeric characters long.

Note

The /SYSTEM_NAME qualifier is not supported for Smart Array controllers.

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> SET GLOBALS/EXPAND_PRIORITY=HIGH/REBUILD_
      PRIORITY=HIGH/SYSTEM_NAME="XXX"/READ_CACHE=50/
      WRITE_CACHE=50
```

Example MSA\$UTIL response for SHOW GLOBALS:

```
Controller:  _$1$GGA1002: (DEFAULT)
Global Parameters:
System Name:  ITA8.2-1
Rebuild Priority: high
Expand Priority: low
Total Cache: 256MB
25% Read Cache: 64MB
75% Write Cache: 192MB
```

SET UNIT

Modifies the attributes of existing units.

Format

SET UNIT <unit_n> / <qualifiers>

Parameters

unit_n

The unit number can be in the range of 0-31.

Qualifiers

/ADG

Specifies that the RAID type for the existing unit is Advanced Data Guard (ADG). Use only with the /MIGRATE qualifier to migrate from any existing RAID level to ADG.

/CACHE

Specifies whether to use the controller's cache for the unit. Caching is On by default. To disable caching, use the /NOCACHE qualifier to the SET UNIT or ADD UNIT command.

/DISKS

Specifies the disks to be used to form the unit. Multiple disks must be enclosed in parentheses. Use only with the /EXPAND qualifier to expand the disks used by the existing unit, as shown in the following format:

```
/DISKS=(disk-numbers[,...])
```

Note

The /DISK qualifier can only be used with the /EXPAND qualifier.

/EXPAND

Allows the specified logical unit and all units in the disk/disk group to utilize more disks. The /EXPAND qualifier does not increase the size of the logical unit; rather, it adds more disks and increases space on individual disks.

To increase the size of the logical unit use the /EXTEND qualifier after you expand the unit.

During expansion of units in a drive group, the RAID level of certain units might change. For example, if a RAID 1 unit with 2 disks is expanded to 3 disks, then the RAID level must change to RAID 5 because the RAID 1 unit does not support an odd number of disks.

When you expand a unit with other units present on the same set of drives, all units undergo volume expansion. Use the following format:

MSA Utility SET UNIT

```
SET UNIT <unit_n>/EXPAND/DISKS=<disk range>
```

Note

The value for disk range must include both pre-expansion disks and the additional disks, for example, 101, (101,112,314, ...).

Only the /DISK qualifier should be used with the /EXPAND qualifier.

Note

The /EXTEND, /EXPAND, and /MIGRATE qualifiers can be used only with the SET UNIT command to change the attributes of an existing unit.

/EXTEND

Increases the size of an existing logical unit. To specify a new size for the unit, use the /SIZE qualifier along with the /EXTEND qualifier. The size specified must be greater than the current unit size. When extending a unit with other units present on the same set of drives, some units might be moved (that is, undergo volume expansion) to make space for the additional size required.

Note that even though the SET UNIT/EXTEND command increases the unit size, you complete other command completing successfully, the increased size takes effect only after you complete other steps (such as SET VOLUME/EXTENSION) successfully at the DCL prompt. For more information, see the *HP OpenVMS System Manager's Volume 1: Essentials* manual.

Use the following format:

```
SET UNIT <unit_n> /EXTEND/SIZE=value
```

Note

The /EXTEND, /EXPAND, and /MIGRATE qualifiers can be used only with the SET UNIT command to change the attributes of an existing unit.

/IDENTIFIER

Specifies the unit number to be used by OpenVMS. The value of the identifier is between 0 and 9999.

Note

The /IDENTIFIER qualifier is not required for Smart Array controllers.

/JBOD

Specifies the RAID type to be JBOD. This is synonymous with RAID 0.

/MIGRATE

Migrates the fault tolerance (RAID) level or stripe size, or both, of an existing logical unit. When migrating a unit with other units present on the same set of drives, some units might undergo volume expansion.

```
SET UNIT <unit_n>/MIGRATE [/RAID_LEVEL=R] [/STRIPE_SIZE=S]
```

where *unit_n* = (0-31), R=(0,1,5), and *S* =(8,16,32,64,128,256)

Note

The /ADG or /JBOD qualifier can also be used instead of the /RAID_LEVEL qualifier.

Only RAID level and stripe size can be modified using the /MIGRATE qualifier.

Cannot migrate any RAID units (RAID 1, RAID 5 and so on) that have spare disks to RAID 0 or JBOD units.

Note

The /EXTEND, /EXPAND, and /MIGRATE qualifiers can be used only with the SET UNIT command to change the attributes of an existing unit.

/RAID_LEVEL=level

Specifies the RAID type of the unit. The supported values for this qualifier are 0 (data stripe), 1 (disk memory), and 5 (data striping with striped parity).

Note

The /RAID_LEVEL qualifier can be used only with the /MIGRATE qualifier.

See the following example for the correct format:

```
/RAID_LEVEL=[(0 | 1 | 5 )]
```

/SIZE=size

Specifies the new size of a unit.

Use the following format:

```
/SIZE=size (GB | MB | KB | %)
```

Note

The /SIZE qualifier can only be used with the /EXTEND qualifier.

/SPARE

/SPARE specifies the disks to be designated as the spare disks. Multiple disks must be enclosed in parentheses. Assigning a spare disk to an unit in a drive group will assign the spare disk to all the configured units in the drive group. If an unit is created on a disk group to which a spare disk is assigned, then the spare disk will be configured to the new unit (if it is not a RAID 0 unit). One spare disk can be assigned to multiple drive groups. Ensure that the size of the spare disk is at least equal to the size of the smallest drive in the drive group.

See the following example for the correct format:

```
/SPARE=(disk_number[,...])
```

MSA Utility SET UNIT

/STRIPE_SIZE=size

Specifies the new stripe size for a given RAID volume. Stripe size must be one of the following values: 8, 16, 32, 64, 128, and 256. Raid 5 and ADG are limited to a maximum 64 KB stripes.

Note

The /STRIPE_SIZE qualifier can be used only with the /MIGRATE qualifier.

/VERBOSE

Provides logging that can be interpreted by engineering.

Examples

1. MSA> SET UNIT 0/CACHE

This command modifies unit 0 and enables the use of the array accelerator for unit 0.

2. SET UNIT 0/EXPAND/DISK=(0,1,2,3)

This command increases the number of disks used by unit 0 and all the other units on the disk group.

3. MSA> SET UNIT 0/SPARE=(100,101)

This command specifies that unit 0 and all the other units on the disk group are spares.

4. MSA> SET UNIT 0/MIGRATE/RAID_LEVEL=1/STRIPE_SIZE=64

This command migrates the RAID level and the stripe size to 64 KB.

5. MSA> SET UNIT 0/EXTEND/SIZE=6GB

This command increases the size of unit 0 to 64 KB.

6. MSA> SET UNIT 0/ID=100

This command sets the identifier for unit 0 to 100.

SHOW CONNECTIONS

Format

SHOW CONNECTIONS / <qualifiers>

Displays the worldwide name, connection name, and profile of each host bus adapter (HBA) attached to the controller. If connections between the HBAs and the MSA utility have been given user-defined names, these names are also displayed.

Note

The SHOW CONNECTIONS command is applicable only to the MSA utility and not Smart Array controllers.

Use the SHOW CONNECTIONS command to verify that all connections to the MSA utility are recognized and defined.

Parameters

None

Qualifiers

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> SHOW CONNECTIONS
Connection Name: abc
Host WWNN = 50060b00-001d25b5
Host WWPN = 50060b00-001d25b4
Profile Name = Default
Unit Offset = 0
Controller 1 Port 1 Status = Online
Controller 2 Port 1 Status = Online
```

This command shows the name of a specific connection to display.

SHOW CONTROLLER

Displays information about the specified controller. If the controller name is not provided, then information of all the connected controllers is displayed.

The SHOW CONTROLLER command is synonymous with the SHOW ADAPTER command.

Note

If the context is set to a specific controller and the controller name is not provided, then SHOW CONTROLLER displays information about only the controller to which the context is set.

Format

SHOW CONTROLLER (controller_name:) / <qualifiers>

Parameters

controller_name:

Qualifiers

/BUS

Displays the bus number of the controller.

/BRIEF

Displays limited information about the controller.

/DEFAULT

Displays information about the default controller.

/NODEFAULT

Displays information about all the controllers.

/FULL

Displays extensive information about the controller.

/REVISION

Displays the firmware version of the controller, along with the default information for the controller.

/SUBSYSTEM

Displays the subsystem ID of the controller.

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> SHOW CONTROLLER

A default controller is not set. All matching controllers
displayed

Controller: _$1$GGA201:
MSA1000      (c) COMPAQ      P56350B9IN2021 Software 4.42
Controller Identifier: 201
NODE_ID = 500805f3-0001b660
SCSI_VERSION = SCSI-3
Supported Redundancy Mode: Active/Standby
Current Redundancy mode: Active/Standby
Current Role: Standby
Device Port SCSI address 6
Host Port_1:
    REPORTED PORT_ID = 500805f3-0001b669
    PORT_1_TOPOLOGY = Not available to this program
Cache:
    102 megabyte read cache  154 megabyte write cache
    Cache is GOOD, and Cache is enabled.
    No unflushed data in cache.
Battery:
    Battery is fully charged.

Controller: _$1$GGA300:
MSA CONTROLLER (c) HP      P56350GX3QN152 Software 6.72
Controller Identifier: 300
NODE_ID = 500508b3-00905ed0
SCSI_VERSION = SCSI-3
Supported Redundancy Mode: Asym Active/Active Active/Standby
    Current Redundancy mode: Asymmetrical Active/Active
    Current Role: Active
    Device Port SCSI address 7
Host Port_1:
    REPORTED PORT_ID = 500508b3-00905ed1
    PORT_1_TOPOLOGY = Not available to this program
Cache:
    128 megabyte read cache  128 megabyte write cache
    Cache is GOOD, and Cache is enabled.
    No unflushed data in cache.
Battery:
    Battery is fully charged.

Controller: _$1$GGA301:
MSA CONTROLLER (c) HP      P56350GX3QN0FS Software 6.72
Controller Identifier: 301
NODE_ID = 500508b3-00905ed0
SCSI_VERSION = SCSI-3
Supported Redundancy Mode: Asym Active/Active Active/Standby
    Current Redundancy mode: Asymmetrical Active/Active
    Current Role: Active
    Device Port SCSI address 6
Host Port_1:
    REPORTED PORT_ID = 500508b3-00905ed9
    PORT_1_TOPOLOGY = Not available to this program
Cache:
    128 megabyte read cache  128 megabyte write cache
    Cache is GOOD, and Cache is enabled.
    No unflushed data in cache.
Battery:
    Battery is fully charged.
```

MSA Utility

SHOW CONTROLLER

```
Controller: _$1$GGA1002:
MSA1000 (c) COMPAQ P56350A9IMN06M Software 4.42
Controller Identifier: 1002
NODE_ID = 500805f3-0001b660
SCSI_VERSION = SCSI-3
Supported Redundancy Mode: Active/Standby
Current Redundancy mode: Active/Standby
Current Role: Active
Device Port SCSI address 7
Host Port_1:
REPORTED PORT_ID = 500805f3-0001b661
PORT_1_TOPOLOGY = Not available to this program
Cache:
102 megabyte read cache 154 megabyte write cache
Cache is GOOD, and Cache is enabled.
No unflushed data in cache.
Battery:
Battery is fully charged.

Adapter: _PKD0:
SA6400 (c) HP P57820FDAPHJE7 Software 1.92
SCSI_VERSION = X3.131:1994 (SCSI-2)
Not currently Redundant
Current Role: Active
(No redundant controller installed)
Cache:
128 megabyte read/write cache
Cache is not configured, and Cache is disabled.
No unflushed data in cache.
Battery:
Battery is not fully charged.

Adapter: _PKE0:
SA6400 EM (c) HP P577C0EDAPH3JK Software 1.92
SCSI_VERSION = X3.131:1994 (SCSI-2)
Not currently Redundant
Current Role: Active
(No redundant controller installed)
Cache:
128 megabyte read/write cache
Cache is not configured, and Cache is disabled.
No unflushed data in cache.
Battery:
Battery is not fully charged.
```

This command displays information about all the controllers because a default controller has not been set.

SHOW DISKS

Displays the following information on all available physical disks: disk number, enclosure bus number and ID, disk size, unit created on the disk, and disks assigned as spares.

Note

This command can be used to view information about a particular disk. For the Smart Array P400 series of controllers, the target ID, connector, and bay are also displayed.

Format

SHOW DISKS / <qualifiers>

Parameters

None

Qualifiers

/AVAILABLE

Lists all the disks on which there are no units configured.

/FULL

Displays the drive model and the drive serial number in addition to the disk information.

Note

For the Smart Array P400 series of controllers, the following information is displayed:

Connector Location
Connector
Enclosure
Bay
WWID
Device Type
Disk Capacity
Device Status
Device Vendor ID
Device Product ID
Device Serial Number
Device Firmware Version
Reserved Area (cfg/status information)
Block Size (bytes/sector)
M&P Data Stamped
Last Failure Reason
Physical Disk Flags

/MEMBER

Displays the disks on which the units are configured.

MSA Utility SHOW DISKS

/SPARE

Displays the disks that are configured as spares for logical units.

/VERBOSE

Provides logging that can be interpreted by engineering.

Examples

1. MSA> SHOW DISKS

This command displays information about all available disks.

2. MSA> SHOW DISK/SPARE
MSA> SHOW DISK/AVAILABLE

These commands display information about designated spare disks and all available disks on which no units are configured.

SHOW GLOBALS

Displays the following global parameters for the specified controller: system name, rebuild and expand-priority settings, and read and write cache settings.

Format

SHOW GLOBALS / <qualifiers>

Parameters

None

Qualifiers

/BRIEF

Displays the global parameters information about the controller.

/FULL

Displays the WWID information in addition to the global parameters.

/VERBOSE

Provides logging that can be interpreted by engineering.

Examples

1. MSA> SHOW GLOBALS
Global Parameters:
System Name: ABC
Rebuild Priority: high
Expand Priority: high
Total Cache: 256MB
50% Read Cache: 128 MB
50% Write Cache: 128 MB

This command displays the global parameters for the controller.

2. MSA> SHOW GLOBALS/FULL
MSA> SHOW GLOBALS/BRIEF

These commands display the WWID parameters and the global parameters for the controller.

MSA Utility

SHOW PROFILE

SHOW PROFILE

Displays all the profile names.

Format

SHOW PROFILE [profile_name] / <qualifiers>

Parameters

profile_name
Displays the specified profile name.

Qualifiers

/VERBOSE
Provides logging that can be interpreted by engineering.

Example

```
MSA> SHOW PROFILE
```

This command displays all profile names.

SHOW OTHER_CONTROLLER

Display information about the controller. The state of this controller can be either ACTIVE or STANDBY.

This is not supported on Smart Array controllers.

Format

SHOW OTHER_CONTROLLER

MSA Utility
SHOW THIS_CONTROLLER

SHOW THIS_CONTROLLER

Displays information about the controller to which the context is set. The state of this controller can be either ACTIVE or STANDBY.

This command applies to both Smart Array and MSA1000/MSA1500 series of controllers.

Format

SHOW THIS_CONTROLLER

SHOW UNIT

Displays the following information for the specified unit:

- Unit identifier (user-defined name)
- Unit status
- List of the disks assigned as data disks to the UNIT
- List of the disks assigned as spares to the UNIT
- RAID level
- Unit size

Format

SHOW UNIT

SHOW UNITS

Displays information about all the logical volumes configured on the controller.

Note

From the command-line interface of MSA1000/MSA1500 controllers, if ACLs are enabled for the units with respect to the connections (HBAs to which these controllers are connected), then only those corresponding UNITS are visible to the respective HBAs.

After the ACLs are enabled, if you create any units on these controllers (from either CLI or MSA\$UTIL), then you must set the ACLs explicitly in order for these units to be visible.

Format

SHOW UNITS <unit_n> / <qualifiers>

Parameters

unit_n

The unit number can range from 0 to 31.

Qualifiers

/BRIEF

Displays information about the unit.

/FULL

Displays the physical configuration information of the unit in addition to the unit information.

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> SHOW UNIT 0/FULL
MSA> SHOW UNIT 1/BRIEF
```

These commands display the physical configuration of the unit as well as the unit information.

SHOW VERSION

Displays the MSA\$UTIL version number.

Format

SHOW VERSION / <qualifiers>

Parameters

None

Qualifiers

/FULL

Displays the WWID, firmware version, and controller revision information in addition to the MSA\$UTIL version.

/VERBOSE

Provides logging that can be interpreted by engineering.

Example

```
MSA> SHOW VERSION/FULL
```

This command displays displays the WWID, firmware version, controller revision information, and the MSA\$UTIL version.

START RECOVER

Initiates the REBUILD operation for all the units configured in the disk group sequentially.

Note that "hiding" the insertion of a new disk from the controller can result in the rebuild operation not starting automatically, even if the controller can detect hot-plug events (such as when connected to an MSA30). Removing a disk while the host system is running, and replacing the disk while the host system is powered off, effectively hides the insertion. The START RECOVER command initiates the rebuild operation in such scenarios also.

This command is supported on both OpenVMS Alpha and OpenVMS I64 platforms.

Note

For additional information about rebuilding volumes, see the SCAN ALL command.

Format

START RECOVER

Parameters

None

Qualifiers

None

Product Directories

This chapter contains information about the OpenVMS Version 8.3–1H1 media kit, including:

- Pointers to installation information
- Directory structure of the OpenVMS Version 8.3–1H1 Operating Environment DVD
- OpenVMS Freeware CDs and pointers to related information
- OpenVMS Open Source Tools CD, containing OpenVMS ports of various Open Source Projects not included elsewhere in the OpenVMS kit
- Product licensing
- Documentation

5.1 OpenVMS I64 Operating Environment

This section describes where the OpenVMS I64 operating environment (OE) is located on the DVD and lists the names and locations of all products that ship on the OpenVMS I64 OE DVD.

To install the OpenVMS Version 8.3–1H1 OE, see the *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*. In addition, before you install the OpenVMS I64, review the release notes in this manual and in the *HP OpenVMS Version 8.3 Release Notes* manual.

Also, refer to *Cover Letter for HP OpenVMS Version 8.3-1H1 for Integrity Servers* for late-breaking information not included in this document.

5.1.1 Directories on the OpenVMS I64 Operating Environment DVD

Table 5–1 lists the top-level directories, documentation directories, and OpenVMS documentation file names for OpenVMS Software Product Descriptions (SPDs).

Product Directories

5.1 OpenVMS I64 Operating Environment

Table 5–1 Directory Structure on the OpenVMS I64 Operating Environment DVD

Directory	File Name/Contents
[AVAILMAN_I640301]	Availability Manager Version version 3.1
[.VMSI18N_I640821]	C/C++ I18N
[.CDSA_I64023]	CDSA Version 2.3
[.CSWS_JAVA_I64030]	CSWS_JAVA Version 3.0
[.CSWS_PERL_I64021]	CSWS_PERL Version 2.1
[.CSWS_PHP_I64013]	CSWS_PHP Version 1.3
[.DCE_I64032]	DCE Version 3.2
[.DCPS_026]	DCPS Version 2.6
[.DECNET_PHASE_IV_I640831H1]	DECnet Phase IV Version 8.3–1H1
[.DECNET_PLUS_I640831H1]	DECnet-Plus Version 8.3–1H1 <i>includes FTAM, OSAK and VT</i>
[.DISTNETBEANS_I640111]	Version 1.1
[.DWMOTIF_I64016]	DECwindows Motif Version 1.6
[.ENTERPRISE_DIR_I64055ECO1]	Enterprise Directory Version 5.5ECO1
[.JAVA_I6401502]	Java™ Version 1.5.0-2
[.KERBEROS_I64031]	Kerberos Version 3.1
[.MGMTAGENTS_I64034]	Management Agents Version 3.4
[.NETBEANS_I64036]	NetBeans Version 3.6
[.OVPA_I64040]	OpenView Performance Agent Version 4.0
[.PERL_I640586]	Perl Version 5.8-6
[.PPU_I64080]	Pay Per Use Version 8.0
[.RTR_I64051]	Reliable Transaction Router Version 5.1
[.SWB_I6401713]	Secure Web Browser Version 1.7–13
[.SWS_I640211]	Secure Web Server Version 2.11
[.SOAP_020]	SOAP Toolkit Version 2.0
[.SSL_I64013]	SSL Version 1.3
[.TCPIP_I64056]	TCP/IP Services for OpenVMS Version 5.6
[.TDC_I64023]	Performance Data Collector (formerly, The Data Collector) Version 2.3-1
[.TDC_RT_I64023]	Version 2.3-1
[.UDDI_010A]	UDDI4J Version 1.0A
[.VMSI18N_I640831H1]	OpenVMS Version 8.3–1H1
[.WBEMCIM_I64]	WBEM Services Version 2.6–1
[.WBEM_PROVIDERS]	WBEM Providers Version 1.5-31
[.WSIT_I64012]	Web Services Integration Toolkit Version 1.2
[.XMLC_I64020]	Version 2.0
[.XMLJ_020]	Version 2.0
[.I640831H1.DOCUMENTATION]	Product documentation in .PS.TXT format

5.2 Directories on the OpenVMS Freeware CD

Three OpenVMS Freeware CDs are included in the OpenVMS Version media kit. The Freeware CDs contain a variety of unsupported software tools and utilities, and OpenVMS ports of common open-source packages. These packages can be useful for a variety of technical, engineering, or educational purposes.

For information about each specific freeware package, refer to the appropriate `FREWARE_README.TXT` file. To access the file, insert the appropriate CD into the CD drive and enter the commands, as appropriate to the Freeware volume being mounted. In the MOUNT commands shown, the `ddcu:` specification is the device name of the particular CD or DVD device present on your OpenVMS system.

```
$ MOUNT/OVERRIDE ddcu:  
$ TYPE ddcu: [FREWARE_README.TXT]
```

Once the appropriate CD is mounted, you can view the disk contents and directory structure using standard DCL commands (such as the `DIRECTORY` command), or you can use commands to launch the Freeware menu system for the appropriate volume. Consult the mount commands on the CD label.

For package-specific licensing and related information, consult the individual packages.

Copies of these packages and packages on older Freeware distributions, submission information for future Freeware distributions, and any updated information for this Freeware distribution are all available at:

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

5.3 Open Source Tools CD

The Open Source Tools CD is a collection of Open Source tools ported to OpenVMS by OpenVMS engineering. These open source tools are provided as free software under the terms of the GNU Lesser General Public License. You can redistribute them and modify them under the terms of the GNU Lesser General Public License as published by the Free Software Foundation Version 2.1 of the License.

HP distributes this library in the hope that it will be useful; however, HP provides NO WARRANTY, without even the implied warranty of MERCHANTABILITY or FITNESS FOR A PARTICULAR PURPOSE. For more details, refer to the GNU Lesser General Public License, which is located in the GNV kit directory on the Open Source Tools CD.

The Open Source Tools CD contains the following ports:

- GNV—An open source, GNU-based, UNIX environment for OpenVMS that provides UNIX application developers, system managers, and users a UNIX-style environment. This facilitates development and porting of UNIX software to OpenVMS. (GNU is a UNIX-like operating system that is free software.) GNV provides a UNIX-like shell (command-line interpreter) environment and a C Run-Time Library (CRTL) supplemental library to provide utilities typically found on UNIX systems. The shell used by GNV is bash (Bourne-Again SHell, from GNU, using the POSIX.2 specification). The Open Source Tools CD contains two GNV kits: one for OpenVMS Alpha and one for OpenVMS I64.

Product Directories

5.3 Open Source Tools CD

- IAS (Intel® Itanium® Assembler/Deassembler)—An OpenVMS I64 port of the open source Itanium assembler available from Intel. It can be used to write low-level Itanium assembler code. A few additional features were added to the assembler to make it more usable on OpenVMS I64. Additional information about the new features and how to use the assembler is in the kit.
- STUNNEL—A program that allows you to encrypt arbitrary TCP connections inside an Secure Sockets Layer (SSL) connection from your OpenVMS system to another machine. Stunnel enables you to secure non-SSL aware applications (such as Telnet, IMAP, and LDAP) by providing the encryption and not requiring changes to the original application. Both images and source are provided.

The Open Source Tools CD also contains:

- SSL sources—Included as an optional layered product with OpenVMS Version 8.3-1H1. The CD provides the SSL Version 1.2 for OpenVMS port sources.
- CDRECORD Sources—Included as a part of OpenVMS Version 8.3-1H1. The full source kit used to create the CD record images is provided in the operating system.
- GnuPG (GNU Privacy Guard)—GNU's tool for secure communication and data storage. It can be used to encrypt data and to create digital signatures. GnuPG includes an advanced key management facility. GnuPG is a complete and free replacement for PGP. Because it does not use the patented IDEA algorithm, it can be used without any restrictions. GnuPG is an RFC 2440 (OpenPGP) compliant application.
- CDSA (Common Data Security Architecture) sources—Included as part of OpenVMS Version 8.3-1H1. The CD provides the full source kit used to create CDSA for OpenVMS 2.3.
- Kerberos sources—Included as part of OpenVMS Version 8.3-1H1. The CD provides the full kit used to create Kerberos for OpenVMS Version 2.1.
- GTK+—An open source, free software library for creating graphical user interfaces.
- libIDL—The IDL compiling library, which is an open source, free software library used for creating tress of CORBA Interface Definition Language (IDL) files.
- Freeware tools tar and zip file in the [.000TOOLS] directory.

5.3.1 Directories on the Open Source Tools CD

Table 5-2 lists the contents and directories on the OpenVMS Open Source Tools CD.

Table 5–2 OpenVMS Open Source Tools CD

Product	Directory
CDRECORD Sources	[.CDRECORD_SOURCE]
CDSA Sources	[.CDSA_SOURCE]
Decompression (unzip) and other tools	[.000TOOLS]
GnuPG	[.GNUPG]
GNV for I64 Version 1.6–4	[.GNV_I64]
GTK+	[.GTK]
IAS (Intel Assembler Source)	[.IAS]
Kerberos Sources	[.KERBEROS_SOURCE]
libIDL	[.LIBIDL]
SSL Sources	[.SSL_SOURCE]
Stunnel	[.STUNNEL]

For current or revised listings, see the OpenVMS Open Source Tools web site at:

<http://h71000.www7.hp.com/opensource/opensource.html>

5.4 Product Licensing

The software on the OpenVMS I64 media belongs to HP. Use of the software is authorized only if you have a valid software license from HP for each of the products.

The License Management Facility (LMF) Product Authorization Key (PAK) allows you to access a software product. A PAK must be registered and loaded before you can install the associated software from this CD. To obtain PAKs, contact your HP support representative or your authorized reseller.

Licensing for OpenVMS I64 differs from licensing for OpenVMS Alpha in that the OpenVMS operating system and certain layered products are provided in operating environments: the Foundation Operating Environment (FOE); the Enterprise Operating Environment (EOE); and the Mission Critical Operating Environment (MCOE). The EOE contains the full contents of the FOE plus additional layered products. The MCOE contains the full contents of the EOE plus OpenVMS Clusters and Reliable Transaction Router (RTR).

You can purchase the FOE license and licenses for additional layered products, or you can choose the EOE or MCOE instead, depending on which layered products you want. For more information about the operating environments and their contents, refer to the *HP Operating Environments for OpenVMS Industry Standard 64 Version 8.3–1H1 for Integrity Servers SPD* at:

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

5.5 OpenVMS Version 8.3–1H1 Documentation

OpenVMS Version 8.3–1H1 documentation consists of two new manuals and one revised manual to supplement the current manuals, which were provided with OpenVMS Version 8.3. These manuals provide the new information you need to install and use Version 8.3–1H1.

The following release documentation is included in the OpenVMS Version 8.3–1H1 media kit:

- Hardcopy documents
 - *HP OpenVMS Version 8.3-1H1 for Integrity Servers Upgrade and Installation Manual*
 - *HP OpenVMS Version 8.3-1H1 for Integrity Servers New Features and Release Notes (this manual)*
 - *Cover Letter for HP OpenVMS Version 8.3-1H1 for Integrity Servers*
 - *HP OpenVMS License Management Utility Manual*
- Online documents
 - Online Documentation Library for HP OpenVMS I64 and Microsoft Windows Platforms CD

Not all books in a documentation set are revised for each release. Typically, each release contains a number of new and revised books, with the remainder of the set consisting of books from previous releases. For example, the OpenVMS Version 8.3-1H1 set includes several new and revised books, with the remaining titles being carryovers from previous releases. The books in each set are the most current versions available.

In addition, all the documentation is available from the HP OpenVMS Systems Documentation Web site:

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

A

ACCEPT UNIT command
 in MSA utility, 4-3
ADD UNIT command
 in MSA utility, 4-4
ANALYZE, 3-13
AST delivery
 POSIX, 3-22

B

binary compatibility
 ISV applications, 2-1
BL860c, 2-7
BL90c, 3-19
Booting, 1-10
 delay, 1-11
BOOT_OPTIONS.COM, 1-11

C

CLUE CONFIG command, 2-9
 /MEMORY, 2-9
CPU_POWER_MGMT default, 3-24
C Run-time Library
 corrections, 3-2
CTRL-P, 3-4

D

DECnet for OpenVMS, 1-7
DECnet-Plus for OpenVMS, 1-7
DECwindows Motif
 See HP DECwindows Motif
DECwindows X11 display server
 peripheral connection requirements, 1-11
DELETE UNIT command
 in MSA utility, 4-6
Documentation, 5-5
Documentation corrections, 3-6
 DCL Dictionary, 3-6
 HP OpenVMS System Management Utilities
 Reference Manual, 3-8
 HP TCP/IP Services for OpenVMS Sockets API
 and System Services Programming Manual,
 3-8

Documentation corrections (cont'd)
 OpenVMS System Services, 3-8
 OpenVMS Version 8.3 New Features and
 Documentation Overview, 3-7
 OpenVMS Version 8.3 Release Notes, 3-7
 TCP/IP Services for OpenVMS Management,
 3-9
DVIS_ADAPTER_IDENT, 2-1
DVIS_MOUNTCNT, 2-1

E

encryption
 encrypt/compress issue, 3-6
 input/output file performance, 3-6
 large file fixes, 3-5, 3-6
EXIT command
 in MSA utility, 4-7

F

Fibre Channel, 3-9
 system disk
 boot requirement, 1-11
Fibre Channel Adapter support, 2-2
firmware
 for Integrity servers, 1-8
firmware, EFI, 3-19
FLASH FIRMWARE command
 in MSA utility, 4-8
Freeware CD, 5-3

G

GiCAP, 3-9
Graphics, 1-12
Graphics Console Support, 2-2

H

HELP command
 in MSA utility, 4-9
HP DECwindows Motif
 firmware update, 1-12
 keyboard, 1-12
 startup messages, 1-12

HP SIM, provisioning from, 3-19
HP TCP/IP Services
 flag values, 3-23
HP TCP/IP Services for OpenVMS, 1-7

I

iconv functions, 3-9
installation and upgrade information
 networking options, 1-7
INSTALL utility
 installing resident images, 3-15
Integrity servers
 firmware, 1-8
internationalization data kit, 2-7

L

Layered products, 5-1
LDAP External Authentication, 2-8
Licenses, 5-3
Linker utility
 enhancements, 2-8
 larger I/O segments, 2-8
 optimizing linker-generated code, 2-8
 signature segment flag, 2-9
 unwind data follows code, 2-9
local area network, 3-19
LOCATE command
 in MSA utility, 4-10

M

Media components, 5-1
MSA
 utility, 4-1
MSA utility
 ACCEPT UNIT command, 4-3
 ADD UNIT command, 4-4
 commands, 4-2
 DELETE UNIT command, 4-6
 EXIT command, 4-7
 FLASH FIRMWARE command, 4-8
 HELP command, 4-9
 LOCATE command, 4-10
 privileges required, 4-1
 READ FIRMWARE command, 4-12
 RESET OTHER_CONTROLLER command,
 4-14
 RESET THIS_CONTROLLER command, 4-13
 restrictions, 4-1
 SCAN ALL command, 4-15
 SET CONTROLLER, 4-16
 SET GLOBALS, 4-17
 SET UNIT command, 4-19
 SHOW CONNECTIONS command, 4-23
 SHOW CONTROLLER command, 4-24
 SHOW DISKS command, 4-27

MSA utility (cont'd)
 SHOW GLOBALS command, 4-29
 SHOW OTHER_CONTROLLER command,
 4-31
 SHOW PROFILE command, 4-30
 SHOW THIS_CONTROLLER command, 4-32
 SHOW UNIT command, 4-33
 SHOW UNITS command, 4-34
 SHOW VERSION command, 4-35, 4-36
 usage summary, 4-1
MSA Utility, 4-1
Multipath Fibre Channel disk devices
 managing list of boot devices, 1-10
multiple servers, provisioning, 3-19

N

name length, InfoServer, 3-19
network
 options, 1-7

O

OE DVD, 5-1
Open Source Tools CD, 5-3
OpenVMS I64
 booting from DVD, 1-10
OpenVMS I64 Boot Manager utility
 removing devices, 1-10
 scanning devices, 1-11
Operating Environment
 See OE DVD, 5-1

P

Parameters, 3-1
Product
 directories, 5-1
 licensing, 5-5
provisioning, 3-18

R

READ FIRMWARE command
 in MSA utility, 4-12
Reboot
 automatic, 1-10
Remedial kits
 obtaining, 1-7
RESET OTHER_CONTROLLER command
 in MSA utility, 4-14
RESET THIS_CONTROLLER command
 in MSA utility, 4-13
restriction
 SYSSLD DRIVER, 3-19

S

satellite system, 3–24
SCAN ALL command
 in MSA utility, 4–15
serial port enumeration, 3–20
SET CONTROLLER
 in MSA utility, 4–16
SET GLOBALS
 in MSA utility, 4–17
SET UNIT command
 in MSA utility, 4–19
SHOW CONNECTIONS command
 in MSA utility, 4–23
SHOW CONTROLLER command
 in MSA utility, 4–24
SHOW DISKS command
 in MSA utility, 4–27
SHOW GLOBALS command
 in MSA utility, 4–29
SHOW OTHER_CONTROLLER command
 in MSA utility, 4–31
SHOW PROFILE command
 in MSA utility, 4–30
SHOW THIS_CONTROLLER command
 in MSA utility, 4–32
SHOW UNIT command
 in MSA utility, 4–33
SHOW UNITS command

 in MSA utility, 4–34
SHOW VERSION command
 in MSA utility, 4–35, 4–36
Smart Array P800 support, 2–2
storage and network I/O controllers support, 2–2
StorageWorks 448c support, 2–2
Support policy for software, 1–6
SYSGEN, 3–24
SYSMAN, 3–24
System Dump Analyzer (SDA)
 utility, 2–9
System Event Log (SEL)
 clearing on Integrity servers, 1–7

T

time zones, 2–10
TZ function, 3–20, 3–23

V

validation, 3–19
virtual connect, 3–24
 failover, 3–24
 fibre SCSI, 3–24

W

write messages, 3–22

