

VSI OpenVMS Version V8.4-2 Cover Letter and Release Notes

March 2016

Preface	4
VSI OpenVMS Version 8.4-2 Cover Letter	4
Intended Audience	4
Document Structure	4
New Features	5
1. Backup – Tunable Compression	5
2. BL890c i4 64 core, and 1.5 TB memory support.....	6
3. BL8x0c i4 Systems.....	7
4. OpenVMS V8.4-1H1 Network Device Driver Enables HPE BL8xx LAN on Motherboard (LOM) NIC	8
5. Time Zone (TZ) Files.....	8
6. VSI OpenVMS V8.4-2 Server Support – pre-i2.....	9
7. VSI OpenVMS V8.4-2 Upgrade Path Support	9
Release Notes	10
VSI OpenVMS Operating System	10
1. ACMELOGIN and LOGINPLUS	10
2. ACPI Error and Warning Messages During Boot	11
3. Alignment Faults	11
4. AUTOGEN	11
5. DCL Commands.....	12
6. Documentation	15
7. LAN Failover Set and Upgrade	15
8. LDDRIVER.....	16
9. License Management Changes	16
10. License Management Utility (LMF) May Fail to Load Valid Licenses	17
11. MONITOR Utility: MONITOR IO FREE LIST SIZE field overflow on large memory systems.....	18
12. MONITOR Utility: /RECORD Loses Default File Naming When Subdirectory is Specified	18
13. PCSI: Saved Database	18
14. Shutting Down TCP/IP, SSH or OpenVMS from an SSH Terminal	19
15. SYS\$ANNOUNCE and SYS\$WELCOME Login Messages after Upgrading to VSI OpenVMS.....	20
16. %SYSTEM-F-BADSTACK error after ^Y and STOP from DCL	20
17. USB Device Is Listed As /NOAVAILABLE After OpenVMS System Reboot ..	21
18. Volume Shadowing.....	21
19. VSI OpenVMS V8.4-2 Operating System Upgrades.....	21
Compilers and Layered Products	23
1. ADA.....	23
2. GNV	23
3. HPE C++	23
4. HPE Layered Products for OpenVMS.....	23
5. HPE System Management Home Page.....	24
6. Java 1.6.....	24
7. Open Source Applications.....	24
8. Oracle Rdb.....	25
9. PCSI Kits from 3 rd Parties	25

10.	VSI Availability Manager for 64 bit Windows	26
11.	VSI DECforms V4.2	26
12.	VSI DECnet Plus Copyright Messages	27
13.	VSI DECset Version 12.9-1 for VSI OpenVMS	27
14.	VSI DECwindows and OPCOM Behavior at System Startup	27
15.	VSI Enterprise Directory 5.7	28
16.	VSI I18N Internationalization Data Kit.....	28
17.	VSI Save Set Manager (SSM) for OpenVMS V1.9-1	29
18.	VSI TDMS Installation Verification Procedure (IVP) Correction.....	30
19.	VSI TDMS Licensing Change	30
20.	VSI WBEM Services for OpenVMS (WBEMCIM)	30
System Improvements.....		33
1.	Improvements in Parameter Displays	33
2.	VSI OpenVMS Updated to UEFI 2.3	33
3.	Support for Larger PROCSECTCNT	34
4.	VSI CSWS V2.4 (Secure Web Server for OpenVMS).....	35
5.	VSI CSWS_JAVA.....	36
6.	VSI CSWS_PHP	37
Appendix A - License Management Configuration Considerations for Mixed-Producer OpenVMS Clusters.....		38
	Introduction	38
	Background	38
	Restriction on Cluster LMF Database Configuration.....	38
	Integrity LMF License Types	39
	Using VSI OpenVMS in a Mixed-Producer Cluster.....	40
	VSI V8.4-1H1 LMF Implementation	41
	Concerns and Questions	45

Preface

VSI OpenVMS Version 8.4-2 Cover Letter

VMS Software, Incorporated (VSI) is pleased to introduce the OpenVMS V8.4-2 operating system for HPE Integrity server platforms. Building upon the success of VSI's previous OpenVMS release, OpenVMS V8.4-2 provides support for:

- 64 core support for BL890c i4 (with hyperthreads off)
- UEFI 2.3
- Network boot on Blades
- Support for rx7640 and rx8640
- Tunable BACKUP compression
- Defect repair
- Software component updates:
 - Enterprise Directory 5.7
 - CSWS Secure Web Server, based on Apache 2.4.12
 - CSWS_PHP
 - CSWS_JAVA
 - WBEM
 - Availability Manager for 64b Windows
 - Updated Time Zone definitions

The VSI engineering team continues to work closely with the HPE OpenVMS engineering team to assure consistency of code between HPE and VSI, and to include HPE defect repairs. VSI V8.4-2 contains defect repair through HPE Update Kit U1100. All of the operating environment testing and qualification was performed on HPE Integrity i4, HPE Integrity i2, and pre-i2 servers and associated IO options.

Intended Audience

This document is intended for all users of the VSI OpenVMS Version 8.4-2 operating system. Read this document before you install, upgrade, or use VSI OpenVMS V8.4-2.

Document Structure

This document contains the following sections:

- *New features*: Newly added functionality available to OpenVMS users.
- *Release notes*: release notes introduced in the current release and notes from the previous version that still apply to the new release. A subheading for each release note indicates either the version of origin (for example, *V8.4-1H1*) or the version when an old note was last updated (for example, a note from Version 8.4-1H1 that was revised for V8.4-2 will be labeled with *V8.4-2*).
- *System improvements*: Internal capability upgrades to the OpenVMS operating system.

New Features

1. Backup – Tunable Compression

The ability to compress file data within a save set during a save operation was added to BACKUP in HPE OpenVMS v8.3. This is done with the BACKUP qualifier `/DATA_FORMAT=COMPRESS`. Using compression yields smaller output save set files at the expense of additional CPU time to compress the data. The compression is done using the ZLIB software compression library. Compressing file data with `/DATA_FORMAT=COMPRESS` is software compression, as opposed to `/MEDIA_FORMAT=COMPACTION` which is hardware compression provided by the tape drive. The two types of compression can be used independently, and simultaneously if desired. Media compaction only applies to save sets on the tape; data compression applies to the save set files regardless of where stored.

Starting with OpenVMS V8.4-2, the compression mechanism may be tuned to provide a greater or lesser amount of compression based on the desired behavior of the BACKUP being done. This is accomplished via the new BACKUP qualifier `/ZLIB_LEVEL=n`, where `n` is an integer from 1 (least compression) to 9 (most compression).

For users who are constrained by fixed size output media which is near capacity, using `/ZLIB_LEVEL=9` will compress the data to the maximum extent possible with ZLIB. Each level of compression requires additional CPU time but produces diminishing amounts of data compression, as the algorithm is re-compressing already compressed data.

For users who are constrained by CPU time to complete a BACKUP but desire some level of compression, using `/ZLIB_LEVEL=1` will provide the initial data compression using the least CPU time possible. Even this lowest amount of compression may provide dramatic results depending on the data being compressed, potentially reaching 70% or even 90% output save set file size reduction. Data that is already compressed, such as .ZIP libraries, will seldom provide any additional size reduction.

The elapsed time to create a BACKUP save set depends on whether the BACKUP is CPU bound or I/O bound. Using compression will reduce the amount of I/O required at the cost of additional CPU time. Whether this results in additional elapsed time will depend on the data being saved. Once a BACKUP is running CPU bound, adding additional compression levels will increase the elapsed time required to complete the operation. The acceptable amount of additional time can be tuned using the values from 1 to 9 for `/ZLIB_LEVEL=n`.

To use the tuning capability, `/ZLIB_LEVEL=n` must be specified in addition to `/DATA_FORMAT=COMPRESS`. The compression level integer value is required syntax.

The ZLIB tuning mechanism is only used during save set creation. During a restore operation from a save set, BACKUP will automatically adapt to whether the data is

compressed and by how much. Save sets created using tuned compression are 100% backwards compatible and may be restored by any version of OpenVMS which supports the software compression qualifier `/DATA_FORMAT=COMPRESS`.

The default compression level used by BACKUP if `/ZLIB_LEVEL=n` is not specified equates to ZLIB level 6. This is unchanged from previous BACKUP versions.

2. BL890c i4 64 core, and 1.5 TB memory support

V8.4-2

VSI OpenVMS V8.4-2 supports BL890c i4 configurations with 64 processor cores (Intel HyperThreads disabled), and 1.5 TB of memory. Previous HPE and VSI OpenVMS versions supported up to 32 cores. The 32-core limit required a BL890c i4 system to be partitioned into single or dual Blade nPars. With 64 core support, a BL890c i4 system can be configured as a single 4 Blade nPar.

Note on 64 core support and performance: With any number of cores, large or small, the way an application environment scales in terms of metrics such as throughput and performance depends greatly on the particular demands the environment places on the operating system. The more processes share data and communicate with one another, the more operating system overhead will limit scalability.

Making the best use of many cores is often achieved by isolating, to as great a degree possible, segments of the workload. *Process affinity* and *user-defined capabilities* are two techniques that can be used to achieve a level of process isolation that will help minimize operating system overhead thus making more processing cycles available to the applications. Not all environments are suited for such techniques and increasing the number of cores may not result in an increase in overall application productivity.

LAN Device Memory Limitation – each Blade in a BL890c i4 has two dual port 10GB on-board NICs, often referred to as the LOM (LAN on Motherboard) ports. On a BL890c i4 in a non-Virtual Connect environment, OpenVMS will configure 16 LAN devices to operate the LOMs. On a BL890c i4 in a Virtual Connect environment, OpenVMS will configure as many as 64 LAN devices (Flex NICs). However, LAN devices need memory to operate. For current versions of OpenVMS, that memory comes from non-paged pool, which is limited in size according to system parameters and the amount of physical memory.

The memory used for network traffic on a given device consists of:

- Receive buffers allocated to the device for incoming traffic,
- Receive buffers filled with receive data handed up to applications and not yet returned to the LAN driver for reuse (receive data in flight),
- Transmit buffers given to the device for transmission (transmit data in flight),

- Transmit buffers owned by the LAN driver awaiting return to applications or to non-paged pool (transmit data in flight),
- Buffers allocated and owned by applications used for network traffic, and
- Data structures allocated by the LAN driver to operate the device.

The number of buffers used by the LAN driver can be modified by a LANCP set device command. You would change the default values when you note lost packet or data overrun errors in the LAN driver counters.

For slow devices with limited bandwidth needs, the amount of memory consumed on behalf of a device is small, and many devices can be configured.

For 10 gigabit devices with high bandwidth needs, the transmit and receive data flowing through the device can reach 2.5 gigabytes per second, which is close to 300k buffers consumed per second, for jumbo frames, and 1600k buffers consumed per second for standard size frames. Even if this peak demand is momentary, you can see that a single 10 gigabit device can put extraordinary pressure on non-paged pool. Therefore, the number of 10 gigabit LAN devices that can be used heavily on a system is limited.

The HPE Quickspecs for the 10 gigabit LAN devices recommend 2 gigabytes of memory per LAN device, meaning that total system memory should be sized large enough to allow 2 gigabytes of memory available to applications and LAN drivers for each LAN device (usually a dual port device). Sizing system memory in such a way enables reasonable use of the LAN devices, but still may exceed the demands on non-paged pool because the maximum size of non-paged pool is limited, and non-paged pool is shared by consumers other than the LAN drivers.

VSI OpenVMS recommends a maximum of 8 active 10 gigabit LAN devices. You can configure and use more, according to the number of LOM, mezzanine, and PCIe slots available on a system, however best performance will be obtained when the memory demands of all the LAN devices on the system are not excessive.

3. BL8x0c i4 Systems

V8.4-2

BL8x0c i4 Blade servers are supported as satellite cluster members in VSI OpenVMS V8.4-2.

BL8x0c i4 Blade servers can be installed and upgraded via an InfoServer over the network in VSI OpenVMS V8.4-2

Note: For satellite cluster and InfoServer operation on BL8x0c i4 Blade servers, the LOM firmware on the BL8x0c Blade must be at 4.9.416.12 (released 19-August-2015). Please consult your HPE hardware support contact for Blade firmware upgrade assistance.

4. OpenVMS V8.4-1H1 Network Device Driver Enables HPE BL8xx LAN on Motherboard (LOM) NIC

V8.4-1H1

HPE Integrity server i4 Blade systems include a new 10 Gigabit Ethernet LAN adapter. A new network device driver SYS\$EWBE3.EXE enables this device. The network device driver SYS\$EWBE3.EXE also supports the HPE BLc NC553m FlexFabric adapter (mezzanine card form factor) and the HPE AT111A PCIe 2-port CNA. These IO cards are available as options for the BL8xxc i4 Blades and the rx2800 i4 server.

System	Adapter	Driver	Device Name Prefix
BL8xx ci4	Emulex 10Gbe BE3 LOM	SYS\$EWBE3.EXE	EW
BL8xx ci4	NC553m (mezz)	SYS\$EWBE3.EXE	EB
rx2800 i4	AT111A	SYS\$EWBE3.EXE	EB

The BE3 chip is a dual port device; normally two devices are configured, unless it is enabled for Flex-10 operation and a Flex-capable interconnect module is installed. The BE3 chip is a Converged Network Adapter; however, the OpenVMS driver only enables the NIC function. The only device-specific characteristic that you can change is flow control, which you enable or disable via LANCP. You can change device-independent characteristics as you would for any other LAN device.

5. Time Zone (TZ) Files

V8.4-2

OpenVMS V8.4-2 updates the time zone database to the IANA 2015f release of time zone (TZ) files. These files maintain system time and make it compliant with UTC rules. The following time zone files have been updated:

AFRICA.;1	EUROPE.;1
ANTARTICA.;1	FACTORY.;1
ASIA.;1	NORTHAMERICA.;1
AUSTRALASIA.;1	PACIFICNEW.;1
BACKWARD.;1	SOUTHAMERICA.;1
ETCETERA.;1	SYSTEMV.;1

The only significant change is the North Korean time zone. In August 2015, North Korean officials announced the country would turn its clocks back by 30 minutes to establish a new standard time. The switch occurred on August 15, 2015. Details of this change are documented at the following web page.

<http://www.timeanddate.com/news/time/north-korea-new-time-zone.html>

6. VSI OpenVMS V8.4-2 Server Support – pre-i2

V8.4-2

VMS Software, Inc (VSI) is happy to announce support for OpenVMS V8.4-2 on HPE Integrity pre-i2 servers. Customers may run VSI OpenVMS V8.4-2 on any server supported by HPE OpenVMS I64, but should consult the following tables for formal support status.

OpenVMS V8.4-2 Server Support			
Server Model	VSI OpenVMS V8.4-2	VSI Technical Support	HPE Technical Support
rx2800 i4	Tested	Supported	Supported
BL860c i4	Tested	Supported	Supported
BL870c i4	Tested	Supported	Supported
BL890c i4	Tested	Supported	Supported
rx2800 i2	Tested	Supported	Supported
BL860c i2	Tested	Supported	Supported
BL870c i2	Tested	Supported	Supported
BL890c i2	Tested	Supported	Supported

Pre-i2 Server Support			
Server Model	VSI OpenVMS V8.4-2	VSI Technical Support	HPE Technical Support
rx1600	Tested	Supported	Not Supported
rx2600	Tested	Supported	Not Supported
rx2620	Tested	Supported	Not Supported
rx2660	Tested	Supported	Not Supported
rx3600	Tested	Supported	Not Supported
rx6600	Tested	Supported	Not Supported
rx7640 ¹	Tested	Supported	Not Supported
rx8640 ¹	Tested	Supported	Not Supported
HPE Integrity Superdome	Not Tested	Not Supported	Not Supported

7. VSI OpenVMS V8.4-2 Upgrade Path Support

VSI supports upgrades to VSI OpenVMS V8.4-2 from VSI OpenVMS V8.4-1H1, HPE OpenVMS v8.4 (with U900, U1000, or U1100 applied), HPE OpenVMS v8.3-1H1 and HPE OpenVMS v8.3.

VSI OpenVMS V8.4-2 Upgrade Path Support	
Upgrade Target Version	VSI Technical Support
VSI Version V8.4-1H1	Supported
HPE Version v8.4 U900, U1000, U1100	Supported
HPE Version v8.3-1H1	Supported
HPE Version v8.3	Supported
HPE Version v8.2-1	Not Supported

¹ Newly supported in V8.4-2

Release Notes

VSI OpenVMS Operating System

1. ACMELOGIN and LOGINPLUS

V8.4-2

The current OpenVMS System Disk supplies two variants of LOGINOUT.EXE and SETP0.EXE images:

- The traditional non sys\$acm-enabled variant (UAF LOGIN)
- The sys\$acm-enabled variant (ACME LOGIN) that supports external authentication

Earlier versions of OpenVMS shipped each variant in different kits, called the ACMELOGIN and LOGIN patch kit; later the LOGINPLUS kit was used to switch between them. Starting with VSI OpenVMS V8.4-2 we have simplified the means for installing ACME (external authentication) LOGIN. There is no longer a LOGINPLUS kit to switch between traditional login and ACME-enabled login. Instead, execute the simple command procedure called SYS\$MANAGER:SYS\$LOGIN_SWITCH.COM to switch from traditional login to ACME-enabled login or to revert back to traditional login.

NOTE: If you upgraded an earlier HPE or VSI OpenVMS system to VSI OpenVMS V8.4-2, and if you were previously using ACME-enabled LOGIN, VSI OpenVMS V8.4-2 has already installed the correct files on your system disk.

When sys\$acm-enabled (ACMELOGIN) variant 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.

Since these images use SYS\$ACM, they will use the authentication policies provided by the ACME agents that have been configured on your system.

A production version of an LDAP ACME agent first became available on HPE OpenVMS Alpha or Integrity Version 8.3 and above; it provides standard LDAP authentication for user login and password-change operations using an LDAPv3--compliant directory server.

A production version of a Kerberos ACME agent first became available on HPE OpenVMS Alpha or Integrity Version 8.3 and above; it provides standard Kerberos authentication for user login and password-change operations.

For more information, please see the ACME_DEV_README.TXT file that resides in SYS\$HELP. There are additional ACME-related documents in SYS\$HELP: ACMELDAP_STD_CONFIG_INSTALL.PDF and .TXT.

2. ACPI Error and Warning Messages During Boot

V8.4-2

VSI OpenVMS might display the following message during boot on cell-based machines (for example, rx8640 or rx7640):

```
ACPI Error (utmutex-0430): Mutex [1] is not acquired, cannot release
[20071219]
```

VSI OpenVMS might display the following message during boot on certain systems that have power management enabled (for example, an rx2660 with the latest processors):

```
ACPI Warning (nseval-0250): Excess arguments - method [_OST] needs 3,
found 7 [20080701]
```

These messages can be ignored. They will be addressed in a future release of VSI OpenVMS.

3. Alignment Faults

V8.4-2

In VSI OpenVMS V8.4-2 a number of alignment faults were fixed. Through analysis of system operation, we discovered that one source of recurring alignment faults on a running system was in the Extended File Cache (XFC). The source of the faults was identified and fixed, which will have a small but positive effect on overall system performance.

4. AUTOGEN

V8.4-1H1

The AUTOGEN procedure is written in DCL, which uses 32-bit arithmetic values. With the advent of very large memory systems, these values can easily overflow and cause unpredictable interpretations for AUTOGEN when setting various system parameters. Similarly, some of these values can overflow when using very large disks for page or swap files.

Some of these issues were addressed starting in HPE OpenVMS v8.4. However, there are still several known issues and VSI is investigating a much different approach for a future release. In most cases, as a temporary workaround you can set a specific value in MODPARAMS.DAT for any affected system parameters.

One easily identifiable case causes a looping behavior during system startup after AUTOGEN, because the Bus-Addressable Pool memory (BAP) parameters are incorrectly calculated. The behavior shows the following symptom at boot:

```
%SYSBOOT-E-BADBAP, NPAG_BAP_* parameters do not match memory layout.
%SYSBOOT-I-SETBAP, Setting NPAG_BAP parameters to default and requesting AUTOGEN/reboot.
```

This is followed by an AUTOGEN and reboot which then encounters the same issue. To avoid this issue, add the following line to MODPARAMS.DAT:

```
NPAG_BAP_MAX_PA = -1
```

This instructs AUTOGEN to use the largest amount possible for the BAP physical memory limit, which is the default behavior. If you already had a particular limit for this value, then the problem should not occur.

Another known symptom involves the page, swap or dump file, when the desired file size should be increased. However, AUTOGEN informs the user to create one or more of these files with SYSGEN, but with a /SIZE=0 qualifier. This recommendation should be ignored. The correct value may be displayed at an earlier point in the AGEN\$PARAMS.REPORT file generated by AUTOGEN, and should be substituted as the desired size rather than 0.

V8.4-2

Many more of the 32-bit arithmetic issues were fixed, for memory size and dump file size calculations.

In addition, the V8.4 alternative AUTOGEN procedure, called AUTOGEN_LM.COM, was deleted. AUTOGEN_LM.COM was intended for use on large memory systems to include some 32-bit calculation fixes. The changes in AUTOGEN_LM.COM were merged back into AUTOGEN.COM so the additional procedure is no longer needed.

5. DCL Commands

V8.4-2

BACKUP/LIST Command Inadvertently Lengthened To Exceed 80 Characters

As part of the 2TB device support in HPE OpenVMS v8.4, the default line length for BACKUP/LIST was inadvertently lengthened to exceed 80 characters. In OpenVMS V8.4-2, the format has been adjusted to fit in an 80-character display. Note that file specifications that exceed 80 characters will still be displayed without truncation.

COPY/FTP Command Qualifier /SSL

In VSI OpenVMS V8.4-1H1, the COPY command qualifier /SSL was not recognized.

```
$ COPY/FTP/SSL *.TXT XYZHOST"username password"::/log
%DCL-W-IVQUAL, unrecognized qualifier - check validity,
spelling, and placement
\SSL\
```

HPE TCP/IP services added the /SSL qualifier in TCP/IP V5.7 but the qualifier was never added to the OpenVMS copy command definition file.

This problem has been fixed in OpenVMS V8.4-2.

Create/Password – New Command

OpenVMS V8.4-2 supports a new command, \$CREATE/PASSWORD with an optional qualifier /LENGTH=*n*, where *n* can be between 6 and 20. This command generates a list of English-pronounceable nonsense words with a length within 2 of the specified length. This command uses the same algorithm for password generation that is used by \$SET PASSWORD/GENERATE but does not require you to specify your old password, nor does it change your current password. (Note that you can get a similar list using SET PASSWORD/GENERATE and an invalid old password as long as you are using UAF-based LOGIN. This command is useful for users employing ACME-based LOGIN and who want to generate a list of passwords for purposes other than OpenVMS credentials).

DELETE/TREE Command May Fail To Delete Files

The DELETE/TREE command failed to delete files when the executing process held delete access through an ACL on the file but did not hold control access on the file. In the example below, STUDENT1 had DELETE access to the directory LEVEL5. However, the directory was not deleted because STUDENT1 does not have control access to LEVEL5.DIR. Because STUDENT2 is the owner and has control access, only STUDENT2 could delete the directory.

```
LEVEL5.DIR;1 [STUDENTS,STUDENT2] (RWE,RWE,RWE,E)
(IDENTIFIER=[STUDENTS,STUDENT1],ACCESS=READ+WRITE+EXECUTE+DELETE)
```

This problem has been fixed in OpenVMS V8.4-2.

DIRECTORY Command Displays a File Size In Blocks of Up To 2 TB

In VSI OpenVMS V8.4-2, the DIRECTORY command displays by default a file size in blocks of up to 2 TB. In previous versions, any size whose value

exceeded 7 digits would be displayed as 7 asterisks, unless the /WIDTH=SIZE=n qualifier was used. The default value for /WIDTH=SIZE=n has been changed from 7 to 10.

Note that for large files where the size is displayed in blocks, the display of the added characters of the file size may force the display to wrap to a second line, depending on which qualifiers are used in addition to /SIZE.

DIRECTORY/DATE=MODIFIED Command Can Display Incorrect Value

In VSI OpenVMS V8.4-1H1, the DIRECTORY/DATE=MODIFIED command can display an incorrect value for files on an ODS-5 disk.

On an ODS-5 disk, separate values indicate the time when file attributes and file data have changed. The VSI OpenVMS V8.4-1H1 behavior is incorrect because it displays the *data modified* time instead of either the *data modified* time or the *attributes modified* time, whichever is the most recent.

VSI OpenVMS V8.4-2 has returned to the behavior seen in HPE OpenVMS v8.4 and earlier, and is also consistent with the behavior for files on ODS-2 disks.

In addition, the DIRECTORY/FULL command has changed to accurately reflect the values of the various date fields:

- The **Revised** time is now called **Modified**
- The **Attributes** time is now called **Attr Mod**
- The **Modified** time is now called **Data Mod**

Lexical Function F\$GETDVI and Terabyte-scale disks.

The lexical function F\$GETDVI can return various elements of information about devices. However, when this lexical retrieves information about terabyte-scale disks, some of the returned values will overflow DCL's 32-bit storage structures. The following elements returned by F\$GETDVI will give erroneous results:

EXPSIZE, FREEBLOCKS, MAXBLOCK, VOLSIZE

Values presented will be negative.

NOTE:

- FREEBLOCKS may read correctly when the disk has been used sufficiently to drop the number of free blocks below the 32-bit threshold.
- This limitation applies only to the DCL F\$GETDVI lexical function. The system service SYS\$GETDVI and the Run-Time Library routine LIB\$GETDVI used from compiled or assembled languages do not have this issue.

DCL command SHOW SYSTEM pages column overflow in large memory systems

The SHOW SYSTEM command will display a string of asterisks in the pages column when the displayed process working set exceeds 999999 pages.

This problem will be addressed in a future version of VSI OpenVMS.

SHOW DEVICE Now Displays Free Blocks Correctly

SHOW DEVICE now displays Free Blocks correctly for values up to two terabytes. Prior to this change, SHOW DEVICE would display asterisks for free block values whose length exceeded nine digits.

6. Documentation

VSI OpenVMS V8.4-2 documentation is provided on an ISO 9660 CD to meet the needs of users of multiple formats. Before you install VSI OpenVMS V8.4-2, read the following documents, which are located on the documentation CD:

- VSI OpenVMS V8.4-2 Installation and Upgrade Manual
- VSI OpenVMS License Management Utility Manual
- VSI OpenVMS V8.4-2 Guide to Media
- VSI OpenVMS V8.4-2 Operating Environment Software Product Description
- VSI OpenVMS V8.4-2 Operating System Software Product Description
- End User License Agreement (EULA)

Additional OpenVMS documentation can be found on the VSI website www.vmssoftware.com and on the HPE website <http://h71000.www7.hp.com/doc/spd.html>

7. LAN Failover Set and Upgrade

V8.4-2

When upgrading from HPE OpenVMS v8.3, the LAN device configuration might not be valid after the upgrade reboot. This behavior has been observed on a clustered system whose network connection is a single LAN failover device comprised of all of the Ethernet devices on the system. After the upgrade reboot, the system does not rejoin the cluster because it cannot communicate through the LAN failover device. To work around this issue, remove the LAN failover device and reboot with the network devices not in a failover set. Then re-create the LAN failover device and reboot into the cluster.

This problem will be addressed in a future version of VSI OpenVMS.

8. LDDRIVER

V8.4-2

The LDDRIVER version supplied with VSI OpenVMS V8.4-2 has been updated from V8.2 to V9.7. Version 9.7 includes many new features, including support for virtual tape (LM) devices. For complete information, see the LD web page at:

www.digiater.nl/lldriver.html

The integrated version identifies itself as LDDRIVER V9.8. It contains exactly one change from the network V9.7 kit required for VSI rebranding support. Do not install the network V9.7 or earlier kits from digiater.nl on VSI OpenVMS V8.4-2.

9. License Management Changes

V8.4-1H1, V8.4-2

VSI OpenVMS releases contain several changes in licenses and license management from HPE OpenVMS versions.

- With the notable exception of Operating Environment (OE) Product Authorization Keys (PAKs), all Digital, Compaq, HPE and third-party PAKs should continue to function as they have on previous versions of OpenVMS.
- VSI supports both the Base Operating Environment (BOE; license name *VSI OPENVMS-I64-BOE*) and the High Availability Operating Environment (HAOE; license name *VSI OPENVMS-I64-HAOE*).
- Only OE PAKs with a producer value of "VSI" will load on VSI OpenVMS releases. An attempt to load a non-VSI OE PAK results in an informational message. Note that in a cluster with a common license database and both HPE and VSI versions of OpenVMS registered, the systems running VSI OpenVMS may issue one or more messages similar to the following example during boot. This is not an error.

```
%LICENSE-I-NOTLOADED, HP OPENVMS-I64-MCOE is not a valid license and was not loaded
```

- The default producer for the F\$LICENSE() DCL lexical function has been changed to "VSI".
- The code in F\$LICENSE has been enhanced to allow the specific producer of "DEC" or "HP" to function in a manner consistent with having omitted the producer in previous versions.

10. License Management Utility (LMF) May Fail to Load Valid Licenses

V8.4-1H1, V8.4-2

VMS Software Inc. has observed some unexpected failures to load valid HPE or VSI licenses in OpenVMS cluster environments. The problem occurs only in clusters with nodes running both VSI and HPE versions of the OpenVMS operating system.

The license load failures occur due to incompatible handling of license units among VSI and HPE nodes. The problem can be avoided with modifications to the LMF database. HPE product PAKs must be restricted to load only on cluster members running HPE OpenVMS versions. Similarly, VSI product PAKs must be restricted to load only on cluster members running VSI OpenVMS versions. If the license PAK restrictions are not implemented, the total cluster license units are improperly shared among the cluster members. The following messages are displayed during boot or during a LICENSE LOAD command when the LMF restrictions are not in place:

```
%LICENSE-W-NOLOAD, license was not loaded for <Product Name>
-LICENSE-F-EXCEEDED, attempted usage exceeds active license limits
%LICENSE-I-UNLOADED, <HP|VSI> <Product Name> has been unloaded
```

Any product licensed under an operating environment level agreement (either BOE or HAOE) is not susceptible to this load failure. The problem only affects products that are purchased a la carte.

If you experience this license load failure on your cluster, follow these steps to clear the problem:

1. Update the <Product Name> LMF PAK records to load on the correct cluster members using include or exclude lists. *Appendix A* describes how to implement include and exclude lists.
2. Use the following commands to unload the license on all cluster nodes and then reload the license on all cluster nodes.
 - a. SYSMAN> SET ENVIRONMENT/CLUSTER
 - b. SYSMAN> DO LICENSE UNLOAD /LOG <Product Name>
 - c. SYSMAN> DO LICENSE LOAD /LOG <Product Name>

The following cluster member cross-reference table summarizes the actions required for different cluster configurations. As an example, row 1 shows that no action is necessary when there are no cluster members running an HPE OpenVMS version.

HPE Cluster Member Present?	VSI OpenVMS V8.4-1H1 Present?	VSI OpenVMS V8.4-2 Present?	Action Required
No HPE version	Yes	Yes	None
HPE any version	No	Yes	Need LMF include or exclude lists
HPE any version	Yes	No	Need LMF include or exclude lists; install V8.4-1H1 LMF patch
HPE any version	Yes	Yes	Need LMF include or exclude lists; install V8.4-1H1 LMF patch

See *Appendix A* for more information on configuring LMF in mixed-producer clusters.

11. MONITOR Utility: MONITOR IO FREE LIST SIZE field overflow on large memory systems

V8.4-1H1, V8.4-2

On systems with over 128GB of memory and large memory free list settings, the FREE LIST SIZE of the MONITOR IO command will overflow and a string of asterisks will be displayed.

This problem will be addressed in a future version of VSI OpenVMS.

12. MONITOR Utility: /RECORD Loses Default File Naming When Subdirectory is Specified

V8.4

When specifying MONITOR's /RECORD qualifier with a subdirectory, the default MONITOR naming convention of .DAT was lost.

This problem has been fixed in VSI OpenVMS V8.4-2. Monitor code is updated to use the OpenVMS SYS\$PARSE() service with a default file name to ensure the default naming convention is inherited in all cases.

13. PCSI: Saved Database

V8.4-1H1, V8.4-2

The rebranding procedure employed during the upgrade from HPE OpenVMS versions v8.3, v8.3-1H1, or V8.4 to a VSI OpenVMS version V8.4-1H1 or V8.4-2 creates a snapshot of the existing PCSI database from prior to the upgrade. This is in the container file (LD disk)
 SYS\$SYSDEVICE:[VMS\$COMMON]PCSI\$DATABASE_HISTORY.DSK.

The following commands allow you to extract pre-upgrade history from the file:

```
$ IF .NOT. F$GETDVI("LD:", "EXISTS") THEN @SYS$STARTUP:LD$STARTUP
$ LD CONNECT SYS$SYSDEVICE:[VMS$COMMON]PCSI$DATABASE_HISTORY /SYMBOL
$ LD_DEVICE = "LDA" + LD_UNIT + ":"
$ MOUNT 'LD_DEVICE' PCSIDATABASE
$ DEFINE/USER PCSI$SYSDEVICE 'LD_DEVICE'
$ DEFINE/USER PCSI$SPECIFIC PCSI$SYSDEVICE:[SYS0.]
$ PRODUCT SHOW HISTORY/REMOTE/FULL
$ DISMOUNT 'LD_DEVICE'
$ LD DISCONNECT 'LD_DEVICE'
```

The container disk will look similar to this:

```
Directory LDA3:[VMS$COMMON]

PCSI$PRE_UPGRADE_HISTORY.LIS;1
PCSI$PRE_UPGRADE_PRODUCTS.LIS;1
SYSEXE.DIR;1

Total of 3 files.

Directory LDA3:[VMS$COMMON.SYSEXE]

HP-I64VMS-AVAIL_MAN_BASE-V0804.PCSI$DATABASE;1
HP-I64VMS-C-V0703-18.PCSI$DATABASE;1
...
PCSI$FILE_SYSTEM.PCSI$DATABASE;1
PCSI$PROCESSOR.PCSI$DATABASE;1
PCSI$ROOT.PCSI$DATABASE;1

Total of 42 files.

Grand total of 2 directories, 45 files.
```

The two text files in [VMS\$COMMON] provide easier access to the data most likely to be required, while the actual database provides an auditable trail of the pre-upgrade history.

14. Shutting Down TCP/IP, SSH or OpenVMS from an SSH Terminal

V8.4-1H1, V8.4-2

Do not run the procedure TCPIP\$SSH_SHUTDOWN.COM from a process logged in via SSH. The procedure will terminate the service without warning. If the system can only be accessed via SSH, the system will be left in an unusable state that can only be resolved by a reboot or from external action from another node, if in a cluster.

This behavior will be addressed in a future release.

Note that the OpenVMS shutdown procedure SHUTDOWN.COM correctly notices that an SSH session is in use and allows the system shutdown to proceed

normally. However, verify that there is nothing in the site-specific SYSHUTDOWN.COM procedure that attempts to shutdown TCP/IP, or the procedure will terminate the process prior to completing the system shutdown.

15. SYS\$ANNOUNCE and SYS\$WELCOME Login Messages after Upgrading to VSI OpenVMS

V8.4-1H1, V8.4-2

The system manager may customize the messages that a user sees when logging into OpenVMS. These are controlled by logical names found in SYS\$MANAGER:SYSTARTUP_VMS.COM. SYS\$ANNOUNCE is displayed before the username and password prompts, and SYS\$WELCOME after a successful login. By default, unless modified by the system manager, these contain a message displaying the operating system name and version.

Because VSI OpenVMS has changed the operating system name, systems upgraded from a previous OpenVMS version may display a stale operating system name. All of the SYS\$MANAGER:SY*.COM procedures are left intact during an upgrade, to preserve any changes made by the system manager during the previous version of OpenVMS, so the prior format of the name will continue after an upgrade.

To use the new format of the name, refer to the SYSTARTUP_VMS.TEMPLATE file, which contains the default content for SYSTARTUP_VMS.COM during an initial system installation. Cut and paste the appropriate section to update the message text in your SYSTARTUP_VMS.COM procedure.

The default operating system name and version format for VSI OpenVMS V8.4-2 is:

```
VMS Software, Inc. OpenVMS (TM) IA64 Operating System, V8.4-2
```

16. %SYSTEM-F-BADSTACK error after ^Y and STOP from DCL

V8.4-1H1

After a Control-Y to interrupt an image or procedure from DCL, a STOP command utilized to terminate the prior context may result in the following exception:

```
Improperly handled condition, bad stack or no handler specified.
Signal arguments:   Number = 0000000000000005
                   Name   = 000000000000000C
                               0000000000000004
                               000000007FF8BCAC
                               FFFFFFFF80002750
                               0000000000000012
```

```
[Register dump and Interrupted Frame information]
```

```
%SYSTEM-F-BADSTACK, bad stack encountered during exception dispatch
```

This error is due to incorrect exception handling which leaves stale information on the user stack. The error is mostly benign since the image has been terminated and the process will continue normally afterwards.

Note that by using EXIT instead of STOP, this problem can be avoided when it is ok to allow any exception handling from the interrupted image to occur (which is true for OpenVMS commands which invoke an image).

This problem was introduced in HPE OpenVMS v8.4 Update0600 and will be fixed in a future release of VSI OpenVMS.

17. USB Device Is Listed As /NOAVAILABLE After OpenVMS System Reboot

V8.4-1H1

When a USB device is connected to an HPE Integrity system running VSI OpenVMS, the device is correctly detected and usable. However, if the system is rebooted while the USB device is connected, the device will be set to /NOAVAILABLE after the OpenVMS boot completes.

This will be fixed in a future release of VSI OpenVMS.

18. Volume Shadowing

V8.4-2

HPE has issued the following advisory for Volume Shadowing: On OpenVMS v8.4 for Alpha and Integrity, using the DISMOUNT keyword (for minicopy) in host-based mini-merge (HBMM) policy definitions can result in inconsistent data between shadow set members when minicopy is used. This same behavior is present in VSI OpenVMS V8.4-1H1 and is fixed by the SHADOWING-V0100 patch kit, available from the HPE Support Center. This problem is fixed in VSI OpenVMS V8.4-2.

19. VSI OpenVMS V8.4-2 Operating System Upgrades

V8.4-2

VSI supports upgrades to VSI OpenVMS V8.4-2 from VSI OpenVMS V8.4-1H1, HPE OpenVMS v8.4 (with U1100, U1000, or U900 applied), HPE OpenVMS v8.3-1H1, and HPE OpenVMS v8.3.

- Perform the VSI OpenVMS V8.4-2 Upgrade on Backup System Disk

VSI strongly recommends that you make a copy of your system disk prior to beginning an upgrade, and that you perform the upgrade on the backup copy.

- VSI OpenVMS V8.4-2 Upgrade Procedure May Report "InsDepErr" Error

The upgrade from HPE OpenVMS versions to VSI OpenVMS V8.4-2 includes a step called rebranding. This step creates a new PCSI database with the required changes, allowing the OpenVMS platform to be upgraded correctly. However, some layered products must be installed in a particular sequence. VSI has identified as many of these as possible, and the upgrade procedure will report any that do not install correctly. At the end of rebranding, any products that did not install correctly will be listed, usually as "InsDepErr". No action is needed; however, please report this through your support channels.

Compilers and Layered Products

VSI has rebranded and released new versions of many HPE layered products. Note that while the VSI version of a layered product is a new version, its functionality is the same as the HPE application that it replaces. VSI and HPE have made every effort to assure that the VSI versions are functionally equivalent to the HPE versions, and include all available defect repairs. Please see the *VSI OpenVMS Version 8.4-2 Guide to Media* for product versions and kit directories.

Copyright notices in layered products have been modified from HPE to VSI. However, the text that appears in some help files, display screens, or other areas may still say HPE. This will be modified in a future release.

1. ADA

V8.4-1H1, V8.4-2

The VSI OpenVMS Layered Product distribution media does not contain Ada. If you require Ada, please contact [Adacore.com](http://adacore.com).

2. GNV

V8.4-2

The GNV Project provides a port of many GNU utilities to OpenVMS with the intent of providing a framework for porting Open Source code to OpenVMS. GNV consists of a base kit and several additional update kits that provide bug fixes and additional functionality. The following GNV page provides download links and additional installation instructions <http://sourceforge.net/projects/gnv>

3. HPE C++

V8.4-1H1, V8.4-2

VSI OpenVMS Layered Product distribution media does not contain C++. If you require C++, you must acquire it from HPE.

4. HPE Layered Products for OpenVMS

V8.4-1H1, V8.4-2

Behavior of HPE OpenVMS Layered Products After Upgrade or Installation

If you upgrade an HPE OpenVMS system to a VSI OpenVMS version and you have HPE layered products on the HPE OpenVMS system disk, those products

should continue to operate after the upgrade to the VSI OpenVMS version. However, if you perform a full installation of a VSI OpenVMS version and then attempt to install an HPE layered product, the layered product installation will fail. If you require assistance with an HPE-supplied layered product kit that does not work, contact VSI.

Layered Products Developed by HPE for OpenVMS

VMS Software, Inc is in process of creating VSI versions of the OpenVMS layered products formerly developed by HPE. If you need to run an HPE application on your VSI OpenVMS system that meets the following criteria, please contact your HPE support representative to understand its support status:

- The application is not yet available from VSI
- The HPE application is qualified with HPE OpenVMS V8.4

5. HPE System Management Home Page

V8.4-1H1, V8.4-2

VSI has not updated HPE System Management Home Page (SMH) for Integrity i4 server support. If HPE SMH is installed on your system and you upgrade to a VSI OpenVMS version, many of the generic operating system parameters will display correctly; however some of the i4 CPU and system board fields will not display. This problem will be fixed in a future release of SMH.

6. Java 1.6

V8.4-1H1, V8.4-2

VSI OpenVMS Operating Environment distribution media does not contain Java 1.6. If you require Java, you must acquire the OpenVMS Java kit from HPE.

7. Open Source Applications

Similar to the layered products, VSI has created new installation kits for many of the Open Source applications that HPE has traditionally provided on the operating environment media. The version number of each PCSI kit has been incremented from the HPE version, but the functionality of the VSI-provided Open Source applications is the same as the HPE-supplied version of the applications. Open Source applications are provided without any warranty or support from VSI, except for the Open Source applications bundled with the BOE. For more information about Open Source applications and sources, see <http://h71000.www7.HP.com/opensource/opensource.html>.

8. Oracle Rdb

V8.4-1H1, V8.4-2

If you use Oracle Rdb, contact Oracle to verify that your Rdb version will install and run on HPE Integrity i4 servers. Oracle Rdb performs an install-time and run-time check for the processor type; some Rdb versions fail this check. Oracle 10g and 11g do not have this same check, so nothing special is needed there.

9. PCSI Kits from 3rd Parties

V8.4-1H1, V8.4-2

Third-party PCSI kits designed to be installed on HPE OpenVMS v8.4 may not install on VSI OpenVMS versions, because of dependencies built into the 3rd party kit. VSI has created a command procedure, `SY$UPDATE:VSI$CVT_KIT_DEPENDENCY.COM`, which will modify a kit to update any such dependencies.

Invoke `SY$UPDATE:VSI$CVT_KIT_DEPENDENCY.COM` with or without parameters; the procedure will prompt for the name of the product and its location. If you invoke it with parameters, they are as follows:

- | | |
|----|---|
| P1 | The name of the product (can include wildcards; a list of kits found will be displayed and you can choose the desired kit). |
| P2 | The location of the kit (device + directory, which can include search lists and wildcards). |
| P3 | If this is "NOCONFIRM" and only one kit matching P1 is found at P2, then the confirmation question is skipped. |

The converted kit will be in the same location as the input kit; its version number will be incremented one higher. The dependencies will be changed in the Product Description File (SOFTWARE statements and functions, and APPLY TO statements). The kit will then be repackaged. If the input kit was compressed, then the output kit will also be compressed.

NOTE:

1. When using the repackaged kit to install a product, you must specify `/OPTIONS=NOVALIDATE_KIT` (no manifest file is created).
2. `SY$UPDATE:VSI$CVT_KIT_DEPENDENCY.COM` should be run from the SYSTEM account, or from an account with full privileges.

10. VSI Availability Manager for 64 bit Windows

V8.4-2

VMS Software, Inc. is pleased to announce Availability Manager Version 3.2-1. This release provides support for 64-bit Microsoft Windows 7 and 8.1. With this release, the functionality of Availability Manager v3.1-2 from HPE has been brought forward to current OpenVMS and Microsoft Windows platforms.

Availability Manager is comprised of two components. The first component is a data collector that executes on the OpenVMS server that is being monitored. The data collector is part of OpenVMS and is enabled by the HAOE license. No installation is required.

The second component is the Windows AM application which has been updated to support 64-bit Microsoft Windows 7 and 8.1.

Testing is currently being performed with Microsoft Windows 10. Support for Microsoft Windows 10 will be introduced in a future release of Availability Manager.

11. VSI DECforms V4.2

V8.4-2

VSI DECforms V4.2 contains the following bug fixes that were made after HPE DECforms V4.0 was released. VSI DECforms V4.1 does not contain these fixes because VSI DECforms V4.1 is based on HPE DECforms V4.0.

1. QXCM1000361159 -- Users with a valid DECforms Run-Time license received OPCOM messages indicating that a DECforms license was not found. With this fix, DECforms Run-Time license holders will no longer see this message.
2. QXCM1000823192 and QXCM1000942585-- Excessive alignment faults were noted during certain DECforms operations. With this fix, the performance issues due to the alignment faults no longer exist.
3. QXCM1000393106 -- An access violation occurred in the DECforms Panel Editor when objects are selected, followed by the keyboard sequence PF1 + 0. With this fix, this access violation will not happen.

12. VSI DECnet Plus Copyright Messages

V8.4-1H1

In VSI DECnet Plus 8.4B that shipped as part of VSI OpenVMS V8.4-1H1, NET\$CONFIGURE.COM and NET\$STARTUP.COM printed HPE copyright messages instead of VSI copyright messages.

These messages have been corrected in VSI DECnet Plus 8.4C that ships with VSI OpenVMS V8.4-2.

13. VSI DECset Version 12.9-1 for VSI OpenVMS

V8.4-1H1, V8.4-2

VSI Code Management System (CMS) Version 4.8-1

In a previous version of CMS, the CMS DIFFERENCES command did not generate the correct output if the second file was specified using the /GENERATION qualifier. In that case, the output of the second file was blank.

This problem has been fixed in VSI CMS V4.8-1.

VSI DECset Version 12.9-1

In VSI OpenVMS V8.4-2, the DECwindows interface to DECset applications works as expected. It did not work in the prior two releases of HPE DECset. This applies to the applications CMS, DTM, ENVMGR, LSE, MMS, PCA, and SCA.

VSI Environment Manager (ENVMGR) Version 1.9

When the DECwindows interface to ENVMGR terminates, an access violation occurs, creating the files DECSET\$ENVMGR_ERR.LOG and DECSET\$ENVMGR_OUT.LOG in the SYS\$LOGIN: directory. This problem also exists in HPE DECset V12.8 ECO1. It did not happen in HPE DECset V12.8 ECO2 or ECO3, because the DECwindows interface did not work in those two versions. The access violation does not result in a loss of data.

14. VSI DECwindows and OPCOM Behavior at System Startup

V8.4-1H1, V8.4-2

Modern systems may include graphics hardware or video capability even if the system is configured as a server, not a workstation. By default, if a supported graphics device exists, the DECwindows components required for workstation support will be loaded during system startup. Additionally, the behavior of OPCOM

and the operator console OPA0: will be modified suitably for a workstation environment.

If your system is *not* a DECwindows graphics workstation, add the following command to SYS\$MANAGER:SYLOGICALS.COM:

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

This will prevent unintended DECwindows actions for server-based systems with graphics capability.

To ensure that all components behave as intended, VSI recommends that you reboot the system after changing the DECW\$IGNORE_WORKSTATION logical name.

The behavior of DECwindows startup has also been enhanced in VSI OpenVMS versions. The OPC\$OPA0_ENABLE logical name is now correctly honored for systems which want to explicitly control the behavior of the operator terminal, for either workstation or server-based systems.

The logical names that control OPCOM behavior are documented in the SYLOGICALS.COM procedure, and its equivalent original state in SYLOGICALS.TEMPLATE.

15. VSI Enterprise Directory 5.7

V8.4-2

VSI is pleased to announce the first release of VSI Enterprise Directory 5.7. The functionality of VSI Enterprise Directory 5.7 is identical to HPE Enterprise Directory v5.6 ECO1 with the addition of fixes contained in limited-release patches ECO2 and ECO3. Note that this version includes the ECO3 Poodle fix released by HPE in August 2015.

16. VSI I18N Internationalization Data Kit

V8.4-2

VSI is pleased to announce VSI I18N, the OpenVMS internationalization data kit that provides locale data and iconv code converters for international software, utilities, and VSI C international run-time routines.

The VSI C RTL provides capabilities that allow application developers to create international software. The VSI C RTL obtains information about a language and a

culture by reading this information from locale files. This kit contains all of the supported locale files.

The kit also contains a set of Unicode codeset converters that allows conversion between any supported codeset (including DEC Multinational Character Set and Microsoft Code Page 437) to any Unicode encoding (UCS-2, UCS-4, or UTF-8).

If you are using these C RTL capabilities, you must install the VMSI18N kit on your system. This kit is located on the VSI OpenVMS Operating Environment DVD.

For the location of this kit, see the *VSI OpenVMS Version 8.4-2 Guide to Media*. Prior to installation, review the release notes for the VMSI18N kit.

17. VSI Save Set Manager (SSM) for OpenVMS V1.9-1

V8.4-2

Product Restrictions and Limitations

If you encounter memory allocation errors, please check your PGFLQUO setting in your account and WSMAX in your SYSGEN. These parameters need to be higher depending upon your use of VSI Save Set Manager. Network access may not be specified for either input or output save set file specifications.

Problems fixed in VSI SSM V1.9-1

1. IVP Fails When on Read-Only Device

When the current default directory was on a read-only device, the IVP failed with the error message:

```
%DCL-E-OPENOUT, error opening device:[SSM.KIT]BLUE.FIL; as output
-RMS-E-WLK, device currently write locked
```

The Save Set Manager IVP has failed.

2. VSI SSM Help Not Available to Non-Privileged Users

If a user did not have elevated privileges, asking for HELP within VSI SSM would result in the following error message:

```
$ saveset help
%HELP-E-OPENIN, error opening SYS$COMMON:[SYSHLP]SSM$HELP.HLB; as input
-RMS-E-PRV, insufficient privilege or file protection violation
```

3. File SSM\$HELP.HLP Not Removed During PRODUCT REMOVE

After uninstalling VSI SSM with a PRODUCT REMOVE command, the file SYS\$HELP:SSM\$HELP.HLB remained on the system.

Changes made in VSI SSM V1.9-1

VSI SSM no longer updates the verb SAVESET in the node-specific DCL tables (SYS\$SPECIFIC:[SYSLIB]DCLTABLES.EXE). Previous versions would update the node-specific file if one existed on your system.

18. VSI TDMS Installation Verification Procedure (IVP) Correction

V8.4-2

In a previous release of VSI TDMS V2.1 the Installation Verification Procedure (IVP) for the run-time only version was improperly built, and did not run. This has been corrected for VSI TDMS RTO V2.1-1.

This problem did not exist for the full development version of VSI TDMS V2.1.

For consistency, both VSI TDMS DEV and VSI TDMS RTO have a version change to V2.1-1, even though the VSI TDMS DEV product has not changed since V2.1.

19. VSI TDMS Licensing Change

V8.4-2

In previous OpenVMS releases, the TDMS and TDMS-RT licenses were treated as special Operating Environment (OE) licenses.

With the changes for OE licensing in OpenVMS V8.4-1H1, a system with an HPE TDMS or TDMS license registered in a common license database might receive a message as follows:

```
%LICENSE-I-NOTLOADED, HP TDMS is not a valid license and was not loaded
```

In OpenVMS V8.4-2 this implementation has changed. VSI TDMS licenses are no longer treated specially, so these informational messages should no longer appear.

20. VSI WBEM Services for OpenVMS (WBEMCIM)

V8.4-2

Product Introduction

VSI is pleased to announce its first release of VSI's WBEM Services kit. VSI WBEM Services V3.0-B151019 is a rebuilt version of the HPE V3.0.2 WBEM Services for

OpenVMS kit. The VSI kit will recognize an HPE Integrity i4 system but is not updated to handle new indications or errors emitted from the HPE Integrity i4 hardware. This kit includes the Poodle fixes.

For additional information on the VSI WBEM Services V3.0 kit, please refer to the WBEMCIM release notes included on the product kit.

Release Note

The WBEM services CIMSERVER process may fail to start under certain circumstances. As a result, WBEM produces CIMSERVER.dmp files in the SYS\$SPECIFIC:[wbem_services] directory. When you display the CIMSERVER.LOG file in the SYS\$SPECIFIC:[wbem_services] directory, you will see one of the following patterns:

```

AST Fault:

EMSWrapper ERROR 618672410 1 Unable to open file EVENTS_VALID_TIME_INTERVAL.TXT
%SYSTEM-F-ASTFLT, AST fault, SP=01E168B0, param=7FFEFFC8, PC=FFFFFFFF8449F6F0,
PS=0000001B, target PC or FV=FFFFFFFFA88FC110, PS=0000001B

Improperly handled condition, image exit forced by last chance handler.
Signal arguments:   Number = 0000000000000007
                   Name   = 000000000000040C
                   0000000001E168B0
                   000000007FFEFFC8
                   FFFFFFFF8449F6F0
                   000000000000001B
                   FFFFFFFFA88FC110
                   000000000000001B

Register dump:

Access Violation:

EMSWrapper ERROR 664805770 1 Unable to open file EVENTS_VALID_TIME_INTERVAL.TXT
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual
address=0000000001E13EF0, PC=FFFFFFFF801F9F00, PS=0000001B

Improperly handled condition, image exit forced by last chance handler.
Signal arguments:   Number = 0000000000000005
                   Name   = 000000000000000C
                   0000000000000000
                   0000000001E13EF0
                   FFFFFFFF801F9F00
                   000000000000001B

Register dump:

```

The VSI WBEMCIM V3.0-B151019 kit is built with the latest C++ compiler to take advantage of C++ improvements, and also includes bug fixes and enhancements. These changes in WBEM result in larger user stack requirements for certain types of requests made to the CIMSERVER process. When a CIMSERVER user thread processes one of the requests, the allocated user stack is exceeded, resulting in a process crash.

To prevent the server process crashes, increase the user stack space allocated to WBEM services worker threads. Until the WBEM Services Product is updated to handle the problem internally, a logical name can be defined to increase the stack size. Use the following logical name definition to increase the thread stack size by 50%. VSI testing has verified that this increase is sufficient to avoid the problem.

```
#! Put this definition in SYLOGICALS.COM or in the startup script
#! used to execute WBEM services startup, typically SYSTARTUP_VMS.COM
$ define /system/exec PEGASUS_VMS_THREAD_STACK_MULTIPLIER 3
```

After the system logical is defined execute the WBEM services startup script:

```
$ @SYS$STARTUP:wbem_services$startup.com
```

This problem will be fixed in a future release of VSI WBEMCIM.

System Improvements

This section describes internal capability upgrades that VMS Software, Inc has made to the OpenVMS operating system.

1. Improvements in Parameter Displays

Several OpenVMS display fields, sized for older, smaller capacity systems, are not sufficient to hold values that reflect the current size of system parameters in more modern OpenVMS environments. For example, the SHOW DEVICE command on a 2TB disk used to display the free blocks as a string of asterisks. Similarly, on large memory systems, it is possible to set WSdef, WSquo, and WSextent to values that would, on older OpenVMS versions, not display correctly in the AUTHORIZE SHOW USER command; instead they displayed a string of asterisks instead of the true value.

These display fields have been expanded in VSI OpenVMS. As VSI and our customers find these instances of display anomalies, we are fixing them by increasing the size of the display fields such that the new, large values will be displayed correctly.

2. VSI OpenVMS Updated to UEFI 2.3

OpenVMS has been updated to UEFI 2.3; this means that vms_loader.efi and related EFI tools (vms_show.efi, vms_set.efi, vms_bcfg.efi) have been recoded to use the related set of UEFI libraries. Previously, OpenVMS used Intel's EFI 1.10, circa 2002, over a decade out of date.

UEFI is the platform firmware API (think "console"). This is internal to the OS Loader and EFI applications so it is not directly visible unless you intend to write UEFI applications. However, it is significant for OpenVMS for several reasons.

HPE Integrity system/platform firmware is based on the open-source UEFI specification (aka Tianocore), most likely UEFI V2.0 or newer depending on how current your system and firmware is. Both the open-source UEFI and its predecessor, Intel EFI, are backwards compatible as far as basic system bootstrap functions are concerned but UEFI offers a much richer set of capabilities.

Moving ahead to these newer libraries allows OpenVMS to take advantage of new API features and replace numerous HPE-specific routines with industry standard library calls. For example, OpenVMS previously developed its own routines to handle EFI/ACPI device paths; therefore, new devices required new code. Now, moving ahead a decade picks up full device path support using industry standard libraries.

There are several similar benefits that offer a better user experience. As OpenVMS moves to the x86-64 architecture many more of these advances will be exposed to the console users, such as high resolution color graphics, touch panel and mouse

support, additional boot devices (i.e. USB devices and others), compliance with the Secure Boot specification, greatly expanded network boot options, elimination of the need for OpenVMS boot drivers, console-level event logging, improved crash dump capability, and more. The intent is to consolidate the disjointed set of EFI tools into a single OpenVMS Boot Manager, certainly for x86-64 and most likely for IA64 too.

UEFI and ACPI go hand-in-hand for device support. UEFI builds the ACPI tables that OpenVMS parses to configure device drivers. This is the essence of “Plug-and-Play” which OpenVMS has only minimally supported. Updating to UEFI was the first of several steps to greatly improve the ability to support a wider range of devices, power states, cooling, etc. in the future. ACPI will be updated in a following release.

The OpenVMS OS loader and related tools do not link against any platform firmware, thus they are not functionally tied to a specific platform firmware revision or vendor. Using UEFI will maintain compatibility with industry standard platforms and no longer be limited to HPE hardware.

Features based on UEFI will continue to be rolled out in future OpenVMS releases. This first step brings OpenVMS up to date and opens the door for modernization and that is really the most noteworthy UEFI “feature” for now.

3. Support for Larger PROCSECTCNT

The PROCSECTCNT system parameter places a limit on the number of process sections (sometimes called program sections) that a single OpenVMS process may have. Each section groups together contiguous virtual pages that have similar attributes.

A single program typically has only a few sections, as the similar pages are grouped together by the compiler and linker when possible. Some programs activate additional shared images, which will raise the number of sections required. Generally, the default value for PROCSECTCNT is sufficient for normal program execution.

However, if an image is sufficiently complex or activates a large number of additional shared images, the limit may be exceeded, which will result in an image activation failure and these error messages:

```
%LIB-E-ACTIMAGE, error activating image
-SYSTEM-F-SECTBLFUL, process or global section table is full
```

In such cases, the value of PROCSECTCNT should be raised to allow the image to properly load and execute.

In prior versions of OpenVMS, the maximum value allowed by SYSGEN was 2048. In OpenVMS V8.4-2, this limit has been raised to 4096.

NOTE: Process section data is stored in a table located in system 32-bit address space (called S0/S1), a very limited resource. To avoid overflowing S0/S1 space, VSI recommends that PROCSECTCNT be increased only if necessary and by reasonable size increments of 128 or 256 when possible.

VSI recommends that you use MODPARAMS.DAT to supply the desired PROCSECTCNT value, and run AUTOGEN afterwards. This will let AUTOGEN attempt to balance the use of S0/S1 resources to reduce the likelihood of overflow.

4. VSI CSWS V2.4 (Secure Web Server for OpenVMS)

V8.4-2

VMS Software Inc. (VSI) is pleased to provide a new VSI-supported version of Secure Web Server for OpenVMS. This release, based on Apache HTTP Server Version 2.4-12 from the Apache Software Foundation, represents a significant update from previous versions, providing many new features and numerous enhancements. Enhancements include reduced memory utilization and more flexible configuration, along with a variety of new loadable modules providing new and enhanced functionality in areas such as session management, request filtering, rate limiting, and proxying. Secure Web Server for OpenVMS Version 2.4 also provides improved support for the development of custom loadable modules.

Subset of Known Problems and Restrictions in Version 2.4

- Do not attempt to use SWS Version 2.4 with older SWS optional kits
- Do not use SWS Version 2.4 with the following optional kits. Using these kits together causes a process crash.
 - CSWS_PERL V2.1 or earlier
 - CSWS_PHP V5.2-17A or earlier
 - CSWS_JAVA (any)

VMS Software Inc. is working to provide updated versions of these optional kits that can be used with SWS Version 2.4.

IMPORTANT: For detailed descriptions of additional new features, enhancements, restrictions, and problems in VSI CSWS V2.4, see the release notes located in the product's documentation directory on the VSI OpenVMS distribution media. Not all new features are provided on VSI OpenVMS. You can find more information at http://httpd.apache.org/docs/2.4/new_features_2_4.html.

Apache HTTP Server Documentation

For information about the Apache web server, see the Apache HTTP Server Version 2.4 documentation at <https://httpd.apache.org/docs/2.4/>. After installing SWS, you can view the web server documentation at <http://your.hostname/manual>, where “*your.hostname*” is that server host name (or IP address) and port number applicable to your installation.

5. VSI CSWS_JAVA

V8.4-1H1, V8.4-2

CSWS_JAVA Version 7.0-29B is available on VSI OpenVMS versions. Apache Tomcat 7.0.29B improves on previous releases in the following ways:

- Added security fixes (CVEs) to this release. For the list of security fixes, see the Apache Tomcat website at <http://tomcat.apache.org>.
- Added support for searching roles in JNDI/LDAP using a different value other than the actual DN or username specified.
- Added stuckThreadNames property as a pair for the stuckThreadIds one; add thread IDs to the log messages.
- Added support for the default error page to be defined in web.xml by defining an error page with a nested location element.
- Annotation scanning is now always performed regardless of the version declared in the web.xml file, expect metadata complete is set to true.

CSWS_JAVA includes the following projects:

- Tomcat (Catalina)
- Ant

For more information, see the Apache Tomcat project web site <http://tomcat.apache.org/>.

Tomcat

CSWS_JAVA includes the following Apache Tomcat technologies:

- JavaServer Pages 2.2
- Java Servlet 3.0

Tomcat is the reference implementation for the Java Servlet 3.0 and JavaServer Pages 2.2 technologies. CSWS_JAVA includes the final Apache Tomcat V7.0.29. Tomcat is a servlet container with a JSP environment. A servlet container is a runtime shell that manages and invokes Java servlets on behalf of users. Servlet containers can be standalone, in-process, or out-of-process. CSWS_JAVA includes support for standalone servlet containers and out-of-process servlet containers. Support for in-process servlet containers (JSSI) will be included in a future version of Tomcat. For more information about Tomcat 7, see <http://tomcat.apache.org/tomcat-7.0-doc/index.html>.

Ant

Ant is also included in CSWS_JAVA. Ant is a partial implementation of the Jakarta Ant subproject, and its use is limited to building the included sample web applications and simple user-written web applications for Tomcat.

For more detail on content, installation, software prerequisites, and other information on CSWS_JAVA, see the release notes bundled with the CSWS_JAVA kit.

6. VSI CSWS_PHP

V8.4-1H1, V8.4-2

CSWS_PHP Version 5.2-17A is available for VSI OpenVMS releases. For general information about PHP, see <http://www.php.net/>.

New features included in CSWS_PHP Version 5.2-17A include:

- CSWS_PHP based on PHP 5.2.17; the following changes are included:
 - Fix for critical vulnerability associated to floating point.
 - Updated time zone database to version 2010.5
 - Upgraded bundled PCRE to version 8.02
 - Rewrote `var_export()` to use `smart_str` rather than output buffering, prevents data disclosure if a fatal error occurs
 - Resolved a critical issue, reported as PHP bug #53632
- Security Fixes

The following security fixes are included in this release:

- CVE-2011-1464
- CVE-2011-4885
- CVE-2011-1148
- CVE-2011-1938
- CVE-2011-2202
- CVE-2011-0421
- CVE-2011-1092
- CVE-2011-0708
- CVE-2010-3807
- CVE-2006-7243

For more details on CSWS_PHP content, installation, software pre-requisites, and other information, see the release notes bundled with the CSWS_PHP kit.

Appendix A - License Management Configuration Considerations for Mixed-Producer OpenVMS Clusters

Introduction

The release of OpenVMS Version 8.4-1H1 by VMS Software Incorporated (VSI) created a new era for OpenVMS clusters. OpenVMS clusters now contain OpenVMS Operating System software provided by two distinct suppliers, Hewlett Packard Enterprise (HPE) and VSI. The existence of two suppliers results in some interesting licensing implications for cluster-wide usage of products supplied by both HPE and VSI. This document describes how to configure the OpenVMS License Management Facility in a mixed-supplier cluster.

Background

The OpenVMS License Management Facility (LMF) enforces the HPE OpenVMS and the VSI OpenVMS licensing agreements. LMF requires a Product Authorization Key (PAK) be applied to your OpenVMS system to license software usage. The company that supplies OpenVMS software is termed the *producer*. The company that issues the PAK is termed the *issuer*. For purposes of this document the supplier, producer, and issuer are identical for each PAK and will be called *producer*. The producers of the OpenVMS Operating System can only have values of VSI, DEC or HP (not HPE).

Restriction on Cluster LMF Database Configuration

The LMF configuration in your OpenVMS cluster must use one common LMF\$LICENSE database or multiple copies of an identical LMF\$LICENSE database. Adhere to the documented standards as described in the section 5.2 of the *OpenVMS License Management Facility Manual*:

If you have multiple system disks in an OpenVMS Cluster environment where all the systems can access one of the system disks, put your common License Database on the readable disk. For any systems that boot from a separate system disk, you must redirect LMF to the License Database. Define the logical name LMF\$LICENSE to be the disk where the database exists.

If you have multiple system disks in an OpenVMS Cluster environment where some systems cannot access one of the system disks, you must keep separate identical License Databases. Whenever one database is modified, you must copy it to update the other databases.

Integrity LMF License Types

There are three different flavors of LMF license PAKs sold for the OpenVMS Integrity Platform. The three flavors are described in this section.

Operating Environment PAKs

The OpenVMS operating system software is licensed using Operating Environment (OE) PAKs. An OE PAK enables usage of a bundle of software products including the operating system. OEs are organized into a hierarchy where the number of software products licensed increases as you move up the hierarchy. VSI OpenVMS Operating Environment products are the Base Operating Environment (BOE) and the High Availability Operating Environment (HAOE). The software products bundled into an OE can be viewed on your OpenVMS system with the DCL command SHOW LICENSE/OE/HIER/FULL. VSI LMF uses a per core license model to enforce the Operating Environment usage. The per core license model is described in the next section. The full hierarchy of the VSI and HPE operating environments is included below for your reference.

Operating Environment Hierarchy Highest (5) to Lowest (1)

```

----- Operating Environment -----
Name  Description      Type Level
HAOE  High Availability  H   5
      GWLM
MCOE  Mission Critical   H   4
      RTR-SVR
      VMSCLUSTER
      VMSCLUSTER-CLIENT
EOE   Enterprise        H   3
      AVAIL-MAN
      RMSJNL
      VOLSHAD
BOE   Base              H   2
      DEGRAM
      OMS
FOE   Foundation        H   1
      OPENVMS-I64
      OPENVMS-USER
      DVNETEND
      DW-MOTIF
      UCX
      TDC
      X500-ADMIN-FACILITY
      X500-DIRECTORY-SERVER
      CIFS

```

Per-Socket, Per-Processor and Per-Core PAKs

When HP introduced Integrity servers, OpenVMS LMF implemented a per processor license enforcement model. At that time per processor meant per CPU. The introduction of multicore processors and CPU threads instigated updates in the license model to adapt to industry changes. Per core licenses were implemented by HPE and per socket licenses were implemented by VSI and HPE.

The presently available per socket model enforces usage of Operating Environments and almost all available VSI and HPE Layered Products. Internally the OpenVMS LMF product enforces per socket licenses as per core licenses. See section 3.2 of the *VSI OpenVMS License Management Utility Manual* for details on per core and per socket licenses.

As an example, assume you purchased the Base Operating Environment along with VMSCUSTER, Volume Shadowing and DECset licenses. You would receive four license PAKs from VSI to be applied on your OpenVMS system. The number of units in each of the PAKs reflects the maximum number of cores available in your per socket license agreement. The DCL command `show license/charge` reports how many license units are required for the system, as in this example:

```
$ show license/charge
```

```
OpenVMS I64/LMF Charge Information for node KILT
This is an HP BL890c i4 (2.53GHz/32.0MB), with 64 cores active
This platform supports up to 8 processor socket(s)
Type: PPL, Units Required: 64 (I64 Per Processor)
Type: PCL, Units Required: 64 (I64 Per Core)
```

The command output above shows that the 8 socket BL890c i4 system KILT requires 64 PAK units to comply with a license agreement to use all available cores on the 8 processor sockets.

Activity PAKs

OpenVMS compilers use Activity PAKs that enforce concurrent user access to the licensed product. In the concurrent user model, the units supplied in the LMF PAK represent the maximum number of user processes that can simultaneously use the product. If your C Compiler PAK has 10 units, a maximum of 10 simultaneous C compiles can be active on your system. Section 3.2.2 of the *VSI OpenVMS License Management Utility Manual* contains additional information on Activity licensing.

Using VSI OpenVMS in a Mixed-Producer Cluster

Exception for Operating Environment Products

Please note that the details in this section do not apply to OpenVMS products included in the Operating Environment level agreement. If your cluster runs at the High Availability OE (HAOE) level, all bundled products, such as VMSCUSTER, will not exhibit the behavior described in this section. Products that are not bundled in HAOE, such as DECset, are susceptible to the behavior.

License Combination

In an OpenVMS cluster containing HPE OpenVMS Integrity members and VSI OpenVMS Integrity members, care must be taken to ensure your license units are combined identically as in a single producer cluster. When LMF charges use of a product, it joins the *producer* and *product* names together to determine the software usage charge. Under the covers, HPE LMF treats the producers DEC and HP as identical. The special treatment allows license PAK units for the same product, like VMSCUSTER, to be combined regardless of the producer name when you purchased the license. In this document the term *synonym* is used to refer to the special treatment of DEC, HP and VSI producers as identical. When the producer name HP is mentioned in the document, HP means the producer names HP and DEC.

VSI V8.4-1H1 LMF Implementation

Problematic Cluster License Configuration Description

When VSI shipped OpenVMS Version 8.4-1H1, one goal was to interact seamlessly with the existing licensing model. VSI therefore added VSI to the DEC, HP synonym list to maintain the previous combination concept. Unfortunately HPE's OpenVMS LMF does not consider VSI a synonym. When a HPE OpenVMS system encounters a VSI license at boot, the HPE system loads and consumes the VSI license units and does not combine the units with a matching HPE license. An example will help demonstrate the issue. Assume the following conditions are present:

- You have a 2-node cluster with 1 HPE OpenVMS v8.4 member and 1 VSI OpenVMS V8.4-1H1 member
- Your HPE OpenVMS v8.4 member is a 2-core system and your VSI OpenVMS V8.4-1H1 member is a 4-core system
- Your common LMF database contains a 2-unit HP producer PAK for product VMSCUSTER and a 4-unit VSI producer PAK for product VMSCUSTER

The cluster licensing agreement requires 6 LMF PAK units to authorize use of VMSCUSTER on both member nodes. However at the initial implementation of VSI OpenVMS V8.4-1H1 it is possible that LMF will charge 8 PAK units to license both members. When there are no restrictions on which member can load the VMSCUSTER licenses, each cluster member node loads the HP and VSI licenses. One of the two scenarios below will occur in these conditions:

Scenario 1 HPE member boots first:

- HPE OpenVMS v8.4 member boots with 2 VMSCUSTER units required
 - HP/VMSCUSTER PAK is loaded, 2 of 2 units consumed
 - VSI/VMSCUSTER PAK is loaded as different product from HP/VMSCUSTER, 2 of 4 units are consumed.
 - 2 VSI/VMSCUSTER units remain

- VSI OpenVMS V8.4-1H1 member boots with 4 VMSCUSTER units required
 - HP/VMSCUSTER license is loaded as a synonym of VSI/VMSCUSTER, 0 of 2 units consumed (no units available)
 - VSI/VMSCUSTER license is loaded as a synonym of HP/VMSCUSTER in *non-compliant mode*, 2 units consumed 4 units required.
 - No VMSCUSTER units remain and LMF inappropriately complains that your license agreement is violated

Scenario 2 VSI member boots first:

- VSI OpenVMS V8.4-1H1 member boots with 4 VMSCUSTER units required
 - HP/VMSCUSTER license is loaded as a synonym of VSI/VMSCUSTER, 2 of 2 units consumed
 - VSI/VMSCUSTER license is loaded as a synonym of HP/VMSCUSTER with 2 of 4 units consumed
 - 2 VSI VMSCUSTER units remain
- HPE OpenVMS v8.4 member boots with 2 VMSCUSTER units required
 - HP/VMSCUSTER license fails to load, 0 units available 2 units required
 - LMF NOLOAD and LMF EXCEEDED messages are displayed for VMSCUSTER
 - VSI/VMSCUSTER license loads and consumes 2 of 4 units as a different product from HP/VMSCUSTER
 - No HP units remain, product VMSCUSTER is not licensed. Remember HP LMF does not consider VSI/VMSCUSTER as a valid license to run OpenVMS clustering.

Addressing LMF Usage Issues on Clusters with VSI OpenVMS V8.4-1H1

There are two requirements for clusters with VSI OpenVMS V8.4-1H1 member nodes: limiting where license PAKs can load, and installing the VMS841H11_LMF-V0100 patch kit.

Limiting The Cluster Members That Can Load a License PAK by OS Version

The most critical step to ensure your cluster is fully licensed is to modify the PAKs in your LMF database to load only on operating system versions supplied by the producer in the PAKs producer value (VSI, DEC or HP). The PAK producer is listed in the output of the DCL command LICENSE LIST /FULL <product>. To list only VSI products you can use the command LICENSE LIST/PROD=VSI /FULL <product>. Likewise you can list all HPE products by adding the /PROD=HP qualifier and value.

Here is example output for a license list of the VSI VMSCUSTER product.

```

Issuer:                VSI
Authorization:         1-VSI-20160211-00027
Product Name:         VMSCUSTER
Producer:             VSI
Units:                2
Version:              0.0
Release Date:         (none)
PAK Termination Date: 18-FEB-2017
Options:              IA64, PPL
Hardware ID:

Revision Level:       7
Status:              Active
Command:             MODIFY
Modified by user:    SYSTEM
Modified on:         17-FEB-2016 20:50:34.46
Exclude:             IA18
  
```

LMF allows you to limit license PAK loads to a specified cluster node or list of cluster nodes. The limits are set using the LICENSE MODIFY command with the /INCLUDE or /EXCLUDE qualifiers. The LICENSE MODIFY command is described in the *VSI OpenVMS License Management Utility Manual*. The required LMF modifications are:

1. VSI PAK modifications:
 - a. Modify ALL VSI license PAKs with an /INCLUDE list containing one or more cluster members running VSI OpenVMS V8.4-1H1 or later
 - b. Modify ALL VSI license PAKS with an /EXCLUDE list value containing ALL cluster members running versions prior to VSI OpenVMS V8.4-1H1
2. HP PAK modifications:
 - a. Modify ALL HP license PAKs with an /INCLUDE list containing one or more cluster members running versions prior to VSI OpenVMS V8.4-1H1
 - b. Modify ALL HP license PAKs with an /EXCLUDE list value containing ALL cluster members running VSI OpenVMS V8.4-1H1 or later

If you already are familiar with LMF include and exclude lists, this task may be straightforward for you. If you are unfamiliar with include and exclude lists, here is an example cluster with one HPE member and 2 VSI members:

SYSTEMS		MEMBERS	
NODE	SOFTWARE	STATUS	Cores and Licenses
IA10	VMS V8.4-1H1	MEMBER	2 cores BOE, VMSCUSTER, VOLSHAD
IA21	VMS V8.4-1H1	MEMBER	2 cores BOE, VMSCUSTER
IA18	VMS V8.4	MEMBER	2 cores FOE, VMSCUSTER, VOLSHAD

Here are sample license modify commands to limit the VMSCLUSTER and VOLSHAD product loads on the above fictitious OpenVMS cluster. The names IA10, IA21 and IA18 are valid only for this example:

```
$ LICENSE MODIFY VMSCLUSTER/PRODUCER=VSI/EXCLUDE=(IA18)
$ LICENSE MODIFY VMSCLUSTER/PRODUCER=HP/INCLUDE=(IA18)
$ LICENSE MODIFY VOLSHAD/PRODUCER=VSI/INCLUDE=(IA10)
$ LICENSE MODIFY VOLSHAD/PRODUCER=HP/EXCLUDE=(IA10,IA21)
```

Installing the VMS841H1I_LMF-V0100 patch kit

VSI requires that you install the VSI OpenVMS patch kit *VMS841H1I_LMF-V0100* kit to restore the pre-V8.4-1H1 LMF license loading behavior.

On VSI OpenVMS V8.4-1H1 systems with all of the LMF PAKS set to load as described in the previous section, you may see confusing LMF load messages during boot or during a LICENSE LOAD command. You may never see these messages on your cluster. The pattern of the confusing messages will look like this:

```
%LICENSE-W-NOLOAD, license was not loaded for <Product name>
-LICENSE-F-EXCEEDED, attempted usage exceeds active license limits
%LICENSE-I-LOADED, HP VMSCLUSTER was successfully loaded with 12
units
```

As displayed the messages read: license load failed, PAK limit exceeded, license loaded successfully! These messages occur when you have multiple PAKS for the same producer/product and all of the PAKs cannot be combined into one single LMF load. The rules for LMF PAK combination are complicated and an explanation is beyond the scope of this document. In the messages displayed above there are 2 distinct LMF loads occurring; the first load fails and the second load succeeds. In prior versions the display of the NOLOAD and EXCEEDED messages only occurred after all the LMF loads were attempted.

Summary for V8.4-1H1

To ensure proper usage of LMF PAK units, limit the nodes on which PAKs can load and install the LMF V8.4-1H1 patch as described in this section.

Addressing LMF Usage Issues on Clusters with VSI OpenVMS V8.4-2 or Later

When all of your VSI cluster members in a mixed-producer environment are running VSI OpenVMS V8.4-2 or later, you must follow the guidelines described in the earlier section titled *Limiting the Cluster Members That Can Load a License PAK by OS Version*. This requirement will remain until all HPE OpenVMS versioned members are migrated to VSI OpenVMS. VSI OpenVMS V8.4-2 includes the fix shipped in the *VMS841H1I_LMF-V0100* kit.

Summary

VSI's OpenVMS Operating System is a distinct, separate product from HPE's OpenVMS Operating System. VSI cannot control HPE's future license and product policies nor can HPE control VSI's future license policies and products. While care is taken to make the transition period from HPE OpenVMS to VSI OpenVMS as smooth as possible, diverging business requirements may emerge during the transition.

To operate an OpenVMS cluster during the transition from HPE to VSI, the HPE and VSI license PAKs must be limited to use on the appropriate producer's operating system version.

Concerns and Questions

Running HP Products on VSI OpenVMS Versions

Sometimes a product produced by HPE enabled by a PAK issued by HPE needs to run on VSI OpenVMS. For example you may be using an OpenVMS layered product with no available VSI product version. VSI OpenVMS will load and authorize any HPE product. However, use care when you migrate to VSI OpenVMS. VSI recommends the following:

1. Make a backup copy of the system disk for the OpenVMS system being migrated to VSI
2. Perform an *upgrade* to the VSI version on the target system disk. The VSI upgrade process will retain any HP products not licensed by VSI on your disk
3. Boot the new disk and use your HP product

If you need to *install* an HPE PCSI kit on a VSI OpenVMS version, you may encounter more issues. The HPE PCSI kit installation dependencies may not be updated to allow installation on any OpenVMS version beyond HPE OpenVMS v8.4. Your options are:

1. Contact HPE support and request a new version of the PCSI kit supporting VSI minimum version VSI OpenVMS V8.4-1H1
2. Contact HPE support to verify if your agreement with HPE allows you extract the PCSI description file, update the dependencies to allow install a VSI install, and repackage the kit
3. If you have a support contract with VMS Software Inc. contact the VSI Support Group at l1-support@vmssoftware.com for instructions on running an unsupported tool that will allow the product to install on VSI OpenVMS.

How do I obtain the VMS841H1I_LMF-V0100 kit?

VSI direct support customers can send email to L1-support@vmssoftware.com and obtain access to the kit. Please reference your VSI support contract number in the email. The VSI LMF patch kit will be made available on the HPE support site as soon as possible. Access the HPE support site or contact your HPE support representative to obtain the kit.