

# VSI OpenVMS Alpha Version V8.4-2L1 Cover Letter and Release Notes

---

January 2017

DO-DVACLR-001

## Contents

<b>1 Preface</b> .....	<b>4</b>
1.1 VSI OpenVMS Alpha Version 8.4-2L1 Cover Letter .....	4
1.2 Intended Audience .....	4
1.3 Document Structure .....	4
<b>2 New Features</b> .....	<b>5</b>
2.1 Backup – Tunable Compression .....	5
2.2 OpenSSL Update .....	6
2.3 Time Zone (TZ) Files.....	7
2.4 VSI OpenVMS AlphaServer Systems Support and Product Licensing .....	7
2.5 VSI OpenVMS Alpha V8.4-2L1 Upgrade Path Support .....	8
2.6 VSI VMS Notary V1.2.....	8
<b>3 Release Notes</b> .....	<b>9</b>
3.1 VSI OpenVMS Operating System .....	9
3.1.1 ACMELOGIN and LOGINPLUS.....	9
3.1.2 Alignment Faults .....	10
3.1.3 AUTOGEN .....	10
3.1.4 DCL Commands.....	11
3.1.5 Documentation .....	13
3.1.6 %IOGEN-E-LOADERR, Error Configuring Device SRA0 .....	14
3.1.7 LAN Failover Set and Upgrade .....	14
3.1.8 LDDRIVER.....	14
3.1.9 License Management Changes .....	15
3.1.10 License Management Utility (LMF) May Fail to Load Valid Licenses .....	15
3.1.11 MONITOR Utility: MONITOR IO FREE LIST SIZE field overflow on large memory systems	16
3.1.12 MONITOR Utility: /RECORD Loses Default File Naming When Subdirectory is Specified	16
3.1.13 PCSI: Access Violation During Installation When Booted From the Installation Medium .	17
3.1.14 PCSI: Saved Database .....	17
3.1.15 Shutting Down TCP/IP, SSH or OpenVMS from an SSH Terminal .....	18
3.1.16 SYS\$ANNOUNCE and SYS\$WELCOME Login Messages after Upgrading to VSI OpenVMS	19
3.1.17 %SYSTEM-F-BADSTACK error after ^Y and STOP from DCL.....	19
3.1.18 Total Timeline Tracking Tool (T4) for VSI OpenVMS Alpha .....	20
3.1.19 USB Device Is Listed As /NOAVAILABLE After OpenVMS System Reboot .....	20

3.1.20	Volume Shadowing .....	20
3.1.21	VSI OpenVMS Alpha V8.4-2L1 Patch Kits.....	21
3.1.22	VSI OpenVMS Installation May Repeat Post-Installation Task Notification.....	23
3.1.23	VSI OpenVMS Upgrade Procedure May Report "InsDepErr" Error.....	23
<b>3.2</b>	<b>Compilers and Layered Products .....</b>	<b>23</b>
3.2.1	Availability Manager Data Analyzer/Data Server .....	23
3.2.2	HPE Layered Products for OpenVMS.....	24
3.2.3	Open Source Kits for VSI OpenVMS Alpha .....	24
3.2.4	PCSI Kits from 3 <sup>rd</sup> Parties .....	24
3.2.5	VSI DECwindows and OPCOM Behavior at System Startup .....	25
3.2.6	VSI TCP/IP V5.7ECO5F .....	26
<b>4</b>	<b>System Improvements .....</b>	<b>30</b>
4.1.1	Improvements in Parameter Displays .....	30
4.1.2	Support for Larger PROCSECTCNT.....	30
<b>5</b>	<b>Guide to VSI OpenVMS Alpha V8.4-2L1 Contents .....</b>	<b>32</b>
5.1.1	VSI OpenVMS Alpha Version 8.4-2L1 Operating System CD Image Directories.....	32
<b>Appendix A License Management Configuration Considerations for Mixed-Producer OpenVMS Clusters.....</b>		<b>33</b>
<b>A.1 Introduction .....</b>		<b>33</b>
<b>A.2 Background .....</b>		<b>33</b>
<b>A.3 Restriction on Cluster LMF Database Configuration.....</b>		<b>33</b>
<b>A.4 Alpha LMF License Types .....</b>		<b>34</b>
<b>A.5 Integrity LMF License Types .....</b>		<b>34</b>
<b>A.6 Using VSI OpenVMS in a Mixed-Producer Cluster.....</b>		<b>36</b>
<b>A.7 VSI LMF Implementation (OpenVMS V8.4-1H1 onward) .....</b>		<b>37</b>
<b>A.8 Concerns and Questions .....</b>		<b>40</b>

# 1 Preface

## 1.1 VSI OpenVMS Alpha Version 8.4-2L1 Cover Letter

VMS Software, Incorporated (VSI) is pleased to introduce the OpenVMS Alpha V8.4-2L1 operating system. This OpenVMS version brings many of the benefits of VSI OpenVMS Integrity V8.4-2 and V8.4-2L1 to Alpha hardware users. In VSI OpenVMS Alpha V8.4-2L1, VSI has rebranded all products in the OPENVMS platform.

VSI OpenVMS Alpha Version 8.4-2L1 is solely developed and marketed by VMS Software, Inc. Hewlett Packard Enterprise (HPE) will not provide support and does not warranty VSI OpenVMS Alpha Version 8.4-2L1.

VSI OpenVMS Alpha V8.4-2L1 contains defect repair through HPE v8.4 Update Kit U1100.

## 1.2 Intended Audience

This document is intended for all users of the VSI OpenVMS Alpha Version 8.4-2L1 operating system. Read this document before you upgrade or use VSI OpenVMS Alpha V8.4-2L1.

## 1.3 Document Structure

This document contains the following sections:

- *New features*: Newly added functionality available to OpenVMS users.
- *Release notes*: Software and documentation problems, restrictions, and corrections.
- *System improvements*: Internal capability upgrades to the OpenVMS operating system.
- *Guide to VSI OpenVMS Alpha V8.4-2L1 Contents*.

## 2 New Features

### 2.1 Backup – Tunable Compression

#### *Alpha V8.4-2L1*

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.

## 2.2 OpenSSL Update

### *Alpha V8.4-2L1*

OpenVMS Alpha V8.4-2L1 introduces VSI SSL1 V1.0, a new release of OpenSSL for OpenVMS. All OpenVMS components in this release that are dependent on OpenSSL have been modified to make use of the new SSL1 offering.

HPE OpenVMS's previous version of SSL (SSL V1.4) is based on the [OpenSSL.org](https://www.openssl.org) code base 0.9.8, which is no longer supported by the OpenSSL community. Many commercial applications and operating systems reject communication to OpenSSL 0.9.8 based targets. Deficiencies in the OpenSSL 0.9.8 feature can be addressed by updating to the OpenSSL V1.0.2 code base.

OpenSSL is used by many operating system functions, networking products, OpenVMS layered products, and open source applications. The prevalence of usage makes OpenSSL a default installation option on OpenVMS systems. OpenVMS Alpha V8.4-2L1 is a coordinated release of OpenSSL V1.0.2 and all software components that use it. This delivery strategy simplifies release packaging and testing and avoids the possibility of complicated patch dependencies.

Note that VSI SSL V1.4 remains unchanged in this release in order to allow existing layered products and customer applications to continue to run. VSI SSL1 V1.0 is designed to co-exist in parallel with VSI SSL V1.4. OpenSSL V1.0.2 is not 100% compatible with V0.9.8, contains new functions and features, and has minor routine interface changes to existing functions. Existing source code may require some minor modifications in order to work with V1.0.2.

VSI OpenVMS Alpha V8.4-2L1 is designed to be compatible with, and allow seamless upgrades from, HPE OpenVMS Alpha releases, including systems that have had the HPE SSL1 set of patches installed.

These SSL-related components are updated in VSI OpenVMS Alpha V8.4-2L1:

- OpenVMS V8.4-2L1 with updates to ENCRYPT, ACME, ACMELDAP
- VSI SSL1 V1.0
- VSI TCP/IP V5.7ECO5F
- VMS Notary

NOTE: Customers are encouraged to migrate their existing SSL based applications to use VSI SSL1 V1.0 as soon as is practicable.

## 2.3 Time Zone (TZ) Files

*Alpha V8.4-2L1*

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>

## 2.4 VSI OpenVMS AlphaServer Systems Support and Product Licensing

*Alpha V8.4-2L1*

### Supported HPE AlphaServer Systems

VSI supports OpenVMS Alpha V8.4-2L1 on the following HPE AlphaServer systems.

Server Model	Tier	Operating System License Maximum Units	Layered Product Maximum Units
DS10, DS15, DS20, DS20E, DS25	Workgroup	75	1050
ES45, ES47	Departmental	700	1100
<i>To be announced</i>	Enterprise	9400	1300

## VSI OpenVMS Alpha Product Licensing

VSI OpenVMS Alpha features a greatly simplified licensing approach. A single license, available with the purchase of a support contract for VSI OpenVMS Alpha, provides the following two PAK groups:

- The ALPHA-SYSTEM PAK group enables the base VSI OpenVMS Alpha operating system.
- The ALPHA-LP PAK group enables all VSI OpenVMS Alpha layered products, some of which are included on the VSI OpenVMS Alpha operating system distribution.

VSI will introduce additional OpenVMS Alpha layered products in the coming months.

## 2.5 VSI OpenVMS Alpha V8.4-2L1 Upgrade Path Support

### *Alpha V8.4-2L1*

VSI supports upgrades to VSI OpenVMS Alpha V8.4-2L1 from previous versions of HPE OpenVMS as noted in the following table. VSI does not support upgrades to VSI OpenVMS Alpha V8.4-2L1 from versions of HPE OpenVMS other than the ones listed here. 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.

<b>VSI OpenVMS Alpha V8.4-2L1 Upgrade Paths: Supported Target Versions</b>
HPE Version v8.4 U900, U1000, U1100, U1200
HPE Version v8.3
HPE Version v7.3-2

**NOTE:** Fresh (from-scratch) installations of VSI OpenVMS Alpha V8.4-2L1 are not supported. Users must upgrade from a valid supported HPE OpenVMS version. Valid upgrade versions are described in the *VSI OpenVMS Alpha Upgrade Manual*, as well as in these Release Notes.

## 2.6 VSI VMS Notary V1.2

### *Alpha V8.4-2L1*

The VMS Notary, part of the OpenVMS operating system, has been updated to use OpenSSL V1.0.2 (SSL1). The VMS Notary allows VSI-signed kits to validate on OpenVMS systems, just as the HPBINARYCHECKER allows HPE-signed kits to validate on OpenVMS.



## 3 Release Notes

### 3.1 VSI OpenVMS Operating System

#### 3.1.1 ACMELOGIN and LOGINPLUS

*Alpha V8.4-2L1*

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. VSI has 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 command procedure SYS\$MANAGER:SYS\$LOGIN\_SWITCH.COM to switch from traditional login to ACME-enabled login or to revert back to traditional login.

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 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 with HPE OpenVMS Alpha Version 8.3; 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.

### 3.1.2 Alignment Faults

*Alpha V8.4-2L1*

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.

### 3.1.3 AUTOGEN

*Alpha V8.4-2L1*

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.

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.

### 3.1.4 DCL Commands

#### *Alpha V8.4-2L1*

#### 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 Alpha V8.4-2L1, 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.

#### Create/Password – New Command

OpenVMS V8.4-2L1 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. (You can get a similar list using SET PASSWORD/GENERATE and an invalid old password as long as you are using UAF-based LOGIN.) The 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 is fixed in OpenVMS Alpha V8.4-2L1.

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

In VSI OpenVMS Alpha V8.4-2L1, 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/FULL Accurately Reflects Values

DIRECTORY/FULL has been modified 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.

## Lexical Function F\$LICENSE Changed to VSI

The F\$LICENSE lexical function has been updated to reflect a change in default producer from DEC to VSI.

### SHOW CLUSTER INCN\_TIME Field May Display Asterisks

SHOW CLUSTER's INCARNATION field is a random 16-digit hexadecimal number that is established when the system is booted. This value differentiates multiple instances of a node with the same SCS identifiers. For example, it may indicate that a node has rebooted. The random value is created by reading the hardware clock; if the clock cannot be read or is suspected of having a bad value, a very large number is used instead, along with some of the clock's low bits.

The INCN\_TIME field is the INCARNATION field formatted as a time value; it can be used to approximate when the system booted or joined a cluster. However, if the hardware clock cannot be successfully read, the field displays asterisks because the "year" which corresponds to the value that was used is approximately 5511, so the value shown as a date will appear very strange.

In this case, the INCARNATION field displays as asterisks in SHOW CLUSTER because SHOW CLUSTER uses 14 digits to display this field; however, the very large number needs 15 hex digits to display correctly. The INCN\_TIME field displays correctly except for the unusual time value far into the future. In such cases, the system continues to operate normally. Likely causes are a faulty hardware clock on the system or a clock battery that needs to be replaced.

### SHOW DEVICE Now Displays Free Blocks Correctly

SHOW DEVICE 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.

### SHOW SYSTEM pages column overflow in large memory systems

SHOW SYSTEM displays 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.

## **3.1.5 Documentation**

### *Alpha V8.4-2L1*

Before upgrading to VSI OpenVMS Alpha V8.4-2L1, read the following documents:

- *VSI OpenVMS Alpha V8.4-2L1 Upgrade Manual*
- *VSI OpenVMS License Management Utility Manual*
- *VSI OpenVMS Alpha V8.4-2L1 Operating System Software Product Description and QuickSpecs*

Additional OpenVMS documentation can be found on the VSI website [www.vmssoftware.com](http://www.vmssoftware.com).

### 3.1.6 %IOGEN-E-LOADERR, Error Configuring Device SRA0

*Alpha V8.4-2L1*

Certain Alpha hardware configurations may issue the following error message during a SYSMAN IO AUTOCONFIGURE command:

```
%IOGEN-E-LOADERR, error configuring device SRA0, driver SYS$SRDRIVER.EXE  
-RMS-E-FNF, file not found
```

The message can be safely ignored. The SRA0 device and corresponding SYS\$SRDRIVER.EXE are Integrity-only devices.

### 3.1.7 LAN Failover Set and Upgrade

*Alpha V8.4-2L1*

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.

### 3.1.8 LDDRIVER

*Alpha V8.4-2L1*

The LDDRIVER version supplied with VSI OpenVMS Alpha V8.4-2L1 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/lddriver.html](http://www.digiater.nl/lddriver.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 Alpha V8.4-2L1.

### 3.1.9 License Management Changes

*Alpha V8.4-2L1*

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

- With the notable exception of Operating System (OS) 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 introduces simplified licensing for the Alpha Operating System with one license for the operating system (license name *VSI ALPHA-SYSTEM*) and another license that enables all Layered Products (license name *VSI ALPHA-LP*).
- 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.

### 3.1.10 License Management Utility (LMF) May Fail to Load Valid Licenses

*Alpha V8.4-2L1*

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 (such as *ALPHA-SYSTEM*) 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 7.3-2, 8.3, or 8.4 Cluster Member Present?	VSI OpenVMS Alpha V8.4-2L1 Present?	Action Required
No	Yes	None
Yes	Yes	Need LMF include or exclude lists

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

### 3.1.11 MONITOR Utility: MONITOR IO FREE LIST SIZE field overflow on large memory systems

*Alpha V8.4-2L1*

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.

### 3.1.12 MONITOR Utility: /RECORD Loses Default File Naming When Subdirectory is Specified

*Alpha V8.4-2L1*

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 Alpha V8.4-2L1. 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.



### 3.1.13 PCSI: Access Violation During Installation When Booted From the Installation Medium

*Alpha V8.4-2L1*

Under certain conditions when installing software while booted from the OpenVMS Installation and Upgrade medium and using Option 3 (Install or upgrade layered products), PCSI may fail with an Access Violation traceback similar to the following:

```
Portion done: 0%...10%...20%...30%...40%...50%...60%...70%
%SYSTEM-F-ACCVIO, access violation, reason mask=00, virtual address=000000000000
0014, PC=00000000009392C4, PS=0000001B
%TRACE-F-TRACEBACK, symbolic stack dump follows
image      module      routine          line      rel PC          abs PC
LBRSHR 0 00000000000312C4 00000000009392C4
PCSI$SHR SPIU_MODULE_LIB_VMS spiu_library_module_execute
20593 0000000000000498 000000000075CAE8
PCSI$SHR SPIU_MODULE_PRIVATE spiu_module_private_main
:
:
```

Use one of the following options to work around this problem:

1. Choose Option 8, Execute DCL commands and procedures

At the \$\$\$ prompt, enter these commands:

```
$$$ ASSIGN/SYSTEM 100 PCSI$DEBUG_FLAGS
$$$ LOGOUT
```

Next, choose option 3 from the installation menu and perform the installation or upgrade again.

2. Perform the product upgrade/installation interactively once the system is booted from the upgraded system disk.

If the Access Violation occurs in this environment, enter the following DCL commands to repeat the PRODUCT INSTALL:

```
$ ASSIGN/SYSTEM 100 PCSI$DEBUG_FLAGS
$ PRODUCT INSTALL...
```

### 3.1.14 PCSI: Saved Database

*Alpha V8.4-2L1*

The rebranding procedure employed during the upgrade from HPE OpenVMS versions v8.3, v8.3-1H1, or v8.4 to VSI OpenVMS version V8.4-2L1 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]

CPQ-AXPVMS-CDSA-V0204-320.PCSI$DATABASE;1
DEC-AXPVMS-DECNET_OSI-V0804.PCSI$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.

### 3.1.15 Shutting Down TCP/IP, SSH or OpenVMS from an SSH Terminal

#### *Alpha V8.4-2L1*

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 is 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.

### 3.1.16 SYS\$ANNOUNCE and SYS\$WELCOME Login Messages after Upgrading to VSI OpenVMS

*Alpha V8.4-2L1*

A 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 Alpha V8.4-2L1 is:

```
OpenVMS (TM) Alpha Operating System, Version V8.4-2L1
```

### 3.1.17 %SYSTEM-F-BADSTACK error after ^Y and STOP from DCL

*Alpha V8.4-2L1*

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.

### **3.1.18 Total Timeline Tracking Tool (T4) for VSI OpenVMS Alpha**

#### *Alpha V8.4-2L1*

The T4 kit that is included with the VSI OpenVMS Alpha distribution is out of date. VSI is currently building new T4 kits, and an updated kit will be available as part of an upcoming VSI OpenVMS Alpha Layered Products release.

If you are using T4 on your existing Alpha system (meaning, before upgrading to the VSI OpenVMS Alpha release) T4 should continue to work with no additional changes after you upgrade to VSI OpenVMS Alpha.

If you are not currently using T4, VSI recommends that you wait until the newer T4 kit is available; do not install the T4 kit included on the VSI OpenVMS Alpha distribution.

### **3.1.19 USB Device Is Listed As /NOAVAILABLE After OpenVMS System Reboot**

#### *Alpha V8.4-2L1*

When a USB device is connected to an HPE server 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.

### **3.1.20 Volume Shadowing**

#### *Alpha V8.4-2L1*

HPE has issued the following advisory for Volume Shadowing: On OpenVMS v8.4, 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 problem is fixed in VSI OpenVMS Alpha V8.4-2L1.

### 3.1.21 VSI OpenVMS Alpha V8.4-2L1 Patch Kits

#### *Alpha V8.4-2L1*

VSI OpenVMS Alpha Version V8.4-2L1 includes code fixes from the following patch kits. Although the “I” in each patch kit name identifies it as an Itanium-based kit, the fixes were applicable to Alpha also and have been included in the Alpha code base:

- VMS842I\_IMGACT-V0100
- VMS842I\_PRCMGT-V0100
- VMS842I\_RMS-V0200
- VMS842I\_VMSINSTAL-V0100

The patch kits address the following problems.

#### **A. VMS842I\_IMGACT-V0100**

**Abstract:** Image activation fails with LOADER-F-NO\_SUCH\_IMAGE

In certain cases the image activator issues an error message for a file that already exists:

```
%DCL-W-ACTIMAGE, error activating image <image-name>
-CLI-E-IMGNAME, image file <full-file-specification>
-LOADER-F-NO_SUCH_IMAGE, the requested image cannot be located
```

#### **B. VMS842I\_PRCMGT-V0100**

**Abstract:** OpenVMS V8.4-2 systems may crash with CWSERR bugcheck

A problem with Special Kernel AST handling within the OpenVMS executive may cause a system crash with a CWSERR bugcheck. Since there is no particular method to determine if a given workload could encounter this issue, VSI recommends that all customers running VSI OpenVMS V8.4-2 systems install this patch kit as a preventative measure.

The bugcheck summary information presented by CLUE CRASH for this problem will have a footprint similar to the following:

```
Bugcheck Type:   CWSERR, Error detected while processing
                  cluster-wide service request
Node:            <nodename>
CPU Type:        <CPU type>
VMS Version:     V8.4-2
Current Process: <process name>
Current Image:   <image name>
Failing PC:      FFFFFFFF.805792D0
                  EXE$NAM_TO_PCB_SCHED_C+00670
Failing PS:      00000000.00000200
Module:          PROCESS_MANAGEMENT
Offset:          000BA7D0
```

### C. VMS842I\_RMS-V0200

**Abstract:** Fix RMS bugchecks when using RMS Global Buffers

A customer site experienced periodic RMS bugchecks with RMS Global Buffers enabled on RIGHTSLIST.DAT and SYSUAF.DAT. This resulted in deletion of the process that incurred the exception, with the following message:

```
%RMS-F-BUG, fatal RMS condition (FFFFFFC0), process deleted
```

This corresponds to the error "BADGBH, Bad Global Buffer Header found" when the RMS Global Buffer is used for opening a file. After this error occurred, the user could no longer log in and a reboot was required to clear the issue.

This fix addresses a small timing window between instructions if a regular file close (or image rundown) is interrupted by a last chance abort rundown.

**Abstract:** Remove RMS executive mode alignment faults

A structure used during RMS name processing is built dynamically on the stack. The base structure definition assumed this would always be quadword aligned, however the stack may only be longword aligned.

By enforcing only longword alignment for this structure, the proper instructions are generated to access it without alignment faults, and with no additional instruction overhead. The routines using this structure are accessed repeatedly during RMS Recovery Unit Journal processing. Other RMS operations may also incur these alignment faults depending on your workload.

This patch kit removes these RMS alignment faults, yielding much improved system performance for these operations.

### D. VMS842I\_VMSINSTAL-V0100

**Abstract:** VMSINSTAL failure on disks with over 1TB of free space

If the system disk free space exceeds 1TB, product installation using the VMSINSTAL mechanism may fail with this message:

```
%CONWRKSSSL-F-NOSYSSPACE, system disk does not contain enough  
free blocks for installation
```

This is caused by limits of 32-bit arithmetic in DCL. The fix corrects the calculation method, allowing installation to very large disks.

### **3.1.22 VSI OpenVMS Installation May Repeat Post-Installation Task Notification**

*Alpha V8.4-2L1*

Products that have multiple dependencies on other products may display required post-installation tasks more than once during kit installation. This happens because PCSI uses a recursive method to ensure that all dependencies are found, but it does not screen previous dependencies under all circumstances. You can safely ignore the duplicated displays; follow the instructions only once. VSI will address this behavior in a future release.

### **3.1.23 VSI OpenVMS Upgrade Procedure May Report "InsDepErr" Error**

*Alpha V8.4-2L1*

The upgrade from HPE OpenVMS versions to VSI OpenVMS Alpha V8.4-2L1 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 VSI support channel.

## **3.2 Compilers and Layered Products**

VSI has rebranded and released new versions of 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 has 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 *Guide to VSI OpenVMS Alpha V8.4-2L1 Contents* 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.

### **3.2.1 Availability Manager Data Analyzer/Data Server**

*Alpha V8.4-2L1*

The Availability Manager Data Analyzer/Data Server kit has traditionally been provided with the OpenVMS operating system distribution kit. However for VSI OpenVMS Alpha V8.4-2L1, the kit will be provided on the Alpha V8.4-2L1 Layered Products distribution kit instead.

### 3.2.2 HPE Layered Products for OpenVMS

*Alpha V8.4-2L1*

#### Behavior of HPE OpenVMS Layered Products After Upgrade

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. 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 VSI 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

### 3.2.3 Open Source Kits for VSI OpenVMS Alpha

*Alpha V8.4-2L1*

VSI has produced several Open Source Tools kits for use with VSI OpenVMS Alpha. For access to any of these kits, please contact us at [info@vmssoftware.com](mailto:info@vmssoftware.com).

Name	Version Number
CIFS (Samba)	1.2ECO1C
CSWS	2.2-1C
CSWS	2.4-3C
CSWS_JAVA	6.0-47A

### 3.2.4 PCSI Kits from 3<sup>rd</sup> Parties

*Alpha V8.4-2L1*

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 3<sup>rd</sup> party kit. VSI has created a command procedure, SYS\$UPDATE:VSI\$CVT\_KIT\_DEPENDENCY.COM, which will modify a kit to update any such dependencies.

Invoke SYS\$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. Run SYS\$UPDATE:VSI\$CVT\_KIT\_DEPENDENCY.COM from the SYSTEM account, or from an account with full privileges.

### **3.2.5 VSI DECwindows and OPCOM Behavior at System Startup**

*Alpha V8.4-2L1*

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 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.

### 3.2.6 VSI TCP/IP V5.7ECO5F

#### Alpha V8.4-2L1

The TCPIP57ECO5F kit (VSI-AXPVMS-TCPIP-V0507-13ECO5F-1) included with VSI OpenVMS Alpha V8.4-2L1 is comprised of a base kit, an ECO kit, and patch kits. After you install the base TCPIP57ECO5F kit, release notes for the component kits will be located in these locations in SYS\$HELP:

Component	Release Note Location
TCPIP 5.7	SYS\$HELP:TCPIP057.RELEASE_NOTES
TCPIP 5.7 ECO5	SYS\$HELP:TCPIP57ECO05.RELEASE_NOTES
TCPIP Telnet Patch	SYS\$HELP:TCPIP_TELNET_PAT57ECO05A.RELEASE_NOTES
TCPIP CVE Patch	SYS\$HELP:TCPIP_CVE_PAT-V57ECO5.RELEASE_NOTES

#### A. Additional Patch Kits Available for VSI TCP/IP: NFS and SSH

Two additional VSI TCP/IP patch kits are included on the VSI OpenVMS Alpha V8.4-2L1 distribution media. The upgrade procedure notifies you that the patches are available.

- VSI AXPVMS TCPIP\_NFS\_PAT V5.7-ECO5A Patch
- VSI AXPVMS TCPIP\_SSH\_PAT V5.7-ECO5D Patch

#### B. VSI TCP/IP\$CONFIG.COM Support for SSH RSA Host Keys

TCPIP\$CONFIG.COM now prompts for the host key type when generating an SSH host key; previously it generated only DSA host keys. Use of an RSA host key allows connectivity with newer SSH client implementations without requiring reconfiguration of the client to support the older DSA host key types.

Here is an example:

```
SSH Configuration
```

```
Service is defined in the SYSUAF.
Service is defined in the TCPIP$SERVICE database.
Service is not enabled.
Service is stopped.
```

```
SSH configuration options:
```

```
1 - Enable service on this node
```

2 - Enable & Start service on this node  
 [E] - Exit SSH configuration

```
Enter configuration option: 1
* Create a new default server host key? [NO]: yes
* Please enter host key type DSA or RSA [RSA]:
  Creating private RSA key file:
TCPIP$SSH_DEVICE:[TCPIP$SSH.SSH2]HOSTKEY
  Creating public RSA key file:
TCPIP$SSH_DEVICE:[TCPIP$SSH.SSH2]HOSTKEY.PUB
```

The SSH CLIENT is enabled.

Do you want to configure SSH CLIENT [NO]:

If an existing key has been in use, SSH clients might refuse connection to the host after you change the key. To correct this, remove the previously used key from the client's list of known keys. Refer to your SSH client documentation for instructions on how to remove keys.

### C. VSI TCP/IP NFS Patch Kit

The VSI TCP/IP V5.7 product includes the NFS kit VSI-AXPVMS-TCPIP\_NFS\_PAT-V0507-ECO5C-4, which is an optional kit that can be selected during the VSI TCP/IP kit installation. This patch kit resolves a problem seen when the /ADF option is used to mount an NFS file system residing on some non-OpenVMS file servers. In this situation, file attributes are not properly written when copying files to the NFS file system. Here is an example of the incorrect behavior:

```
TBD_$ tcpip MOUNT DNFS100: /adf/HOST="10.5.117.2"/PATH="/usr/nfs-test"/system
TBD_$ backup /log test.txt dnfs100:[NFS-TEST]test.txt
%BACKUP-S-CREATED, created DNFS100:[NFS-TEST]TEST.TXT;3
%BACKUP-W-WRITEERR, error writing DNFS100:[NFS-TEST]TEST.TXT;2
-SYSTEM-W-FCPWITERR, file processor write error
%BACKUP-E-WRITEATTR, error writing attributes for DNFS100:[NFS-TEST]TEST.TXT;2
-SYSTEM-W-FCPWITERR, file processor write error
%BACKUP-S-CREATED, created DNFS100:[NFS-TEST]TEST.TXT;2
%BACKUP-W-WRITEERR, error writing DNFS100:[NFS-TEST]TEST.TXT;1
-SYSTEM-W-FCPWITERR, file processor write error
%BACKUP-E-WRITEATTR, error writing attributes for DNFS100:[NFS-TEST]TEST.TXT;1
-SYSTEM-W-FCPWITERR, file processor write error
%BACKUP-S-CREATED, created DNFS100:[NFS-TEST]TEST.TXT;1
```

### D. VSI TCP/IP SSH Patch Kit Addresses Multiple Problems

The VSI TCP/IP V5.7 product includes the SSH patch kit VSI-AXPVMS-TCPIP\_SSH\_PAT-V0507-ECO5D-4 which is an optional kit that can be selected during the VSI TCP/IP kit installation. The patch kit addresses the following problems:

1. When authenticating an LDAP user account using RedHat/Fedora Directory Server, the server may send password expiration warning in the Bind Response using Netscape controls. But for SSH sessions, the respective warning message doesn't get displayed in OpenVMS at login time.
2. ACMELOGIN images, which are bundled with the OpenVMS operating system, are a pre-requisite for LDAP authentication. However, with ACMELOGIN installed, the SET PASSWORD command will not work in SSH sessions, even though it works fine in other terminals like Telnet, DECnet, etc. This restriction is documented in the VSI TCP/IP 5.7 ECO5 release notes, as well as the ACME LDAP Installation Guide.

To overcome this limitation, define the following logical before starting SSH server:

```
$ DEFINE /SYSTEM TCPIP$SSH_SERVER_USE_LOGINOUT 1
```

This prompts SSH server to use [SYSEXE]LOGINOUT.EXE image for authentication (as is done in Telnet and DECnet terminals). Now the "SET PASSWORD" command will work fine in SSH sessions.

Please note that enabling this feature has following side effects:

- a) The LOGINOUT component will display a default Welcome message if SYS\$WELCOME logical is not defined. This is the behavior in Telnet, DECnet, etc. since they use LOGINOUT. On the other hand, SSH server will not display any Welcome message if SYS\$WELCOME logical is not defined. Enabling TCPIP\$SSH\_SERVER\_USE\_LOGINOUT will cause SSH server to behave like Telnet - i.e. display a default Welcome message when SYS\$WELCOME is undefined.
- b) By default, the Audit server logs the LOGIN event for DCL process in SSH session as "Detached process login". But if TCPIP\$SSH\_SERVER\_USE\_LOGINOUT is enabled, the respective event will be "Local interactive login".
- c) In the case of LDAP users, the ACME server will perform user-authentication twice if TCPIP\$SSH\_SERVER\_USE\_LOGINOUT is enabled - the first one initiated by SSH server, and the second one by LOGINOUT image. Hence the timestamp "Last interactive login" displayed at start of interactive session will be almost same as the current time. This is because the timestamp is updated twice in quick succession by the ACME server.
- d) The logical TCPIP\$SSH\_SERVER\_USE\_LOGINOUT doesn't take effect in the following scenarios:
  - SSH sessions using remote command mode
  - SFTP/SCP sessions
  - SSH logins using public-key or host-based authentication
  - OpenVMS users with secondary password
- e) While displaying the Welcome message, width of the terminal may be limited to 80 characters, if LOGINOUT is enabled.

3. For an LDAP user, if the password has already expired, then the SSH authentication simply fails without any indication of password expiry.

Note: The primary reason for this issue is a bug in the ACME LDAP agent. Hence this issue is not exclusive to SSH; it occurs even in Telnet, DECnet, as well as the system console prompt. VMS Security Engineering will release a patch for ACME LDAP agent to address this issue (Elevation: QXCM1001425603 / 4754317355). However, to prevent the problem in SSH sessions, a few related changes are required in the SSH server as well. Hence for SSH users, both ACME LDAP agent patch as well as the SSH server patch (V5.7-ECO5D) must be installed in the system.

4. SSH connections from a client which mandates support for the diffie-hellman-group14-sha1 key exchange method fail.

## 4 System Improvements

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

### 4.1.1 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.

### 4.1.2 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 V8.4-2L1, 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.

## 5 Guide to VSI OpenVMS Alpha V8.4-2L1 Contents

The VSI OpenVMS Alpha V8.4-2L1 release includes the following items:

- VSI OpenVMS Alpha Version 8.4-2L1 Operating System CD image
- *VSI OpenVMS Alpha Version 8.4-2L1 Cover Letter and Release Notes*
- *VSI OpenVMS Alpha Version 8.4-2L1 Upgrade Manual*. (As an aid to your upgrade experience, the Alpha V8.4-2L1 manual includes the *VSI OpenVMS Integrity Version 8.4-2 Installation and Upgrade Manual*)
- *VSI OpenVMS Alpha Version 8.4-2L1 Software Product Description and QuickSpecs*

### 5.1.1 VSI OpenVMS Alpha Version 8.4-2L1 Operating System CD Image Directories

Table 1 lists the names, version numbers, and directories of products found on the VSI OpenVMS Alpha Version 8.4-2L1 Operating System CD image. Before you install VSI OpenVMS Alpha V8.4-2L1, see the *VSI OpenVMS Alpha Version 8.4-2L1 Upgrade Manual*. Review this *VSI OpenVMS Alpha Version 8.4-2L1 Cover Letter and Release Notes* for problems, changes, restrictions, enhancements and new features.

Table 1 Products on VSI OpenVMS Alpha Version 8.4-2L1 OS CD Image

VSI Product Name	Version	Directory
VSI Availability Manager Base	8.4-2L1	[KITS.AVAIL_MAN_BASE_KIT]
VSI CDSA	2.4-320A	[KITS.CDSA]
VSI DECnet Phase IV	8.4-2L1	[KITS.DECNET_PHASE_IV_I640842L1_KIT]
VSI DECnet Plus	8.4-D	[KITS.DECNET_PLUS]
VSI DECwindows Motif	1.7-F	[KITS.DWMOTIF]
VSI DECwindows Motif Support	8.4-2L1	[KITS.DWMOTIF_SUPPORT_I640842L1_KIT]
VSI FTAM	V4.1-B	[KITS.DECNET_PLUS]
VSI HPBINARYCHECKER	1.1-A	[KITS.HPBINARYCHECKER]
VSI Kerberos	3.1-152A	[KITS.KERBEROS]
VSI OSAK	3.0-W	[KITS.DECNET_PLUS]
VSI Secure Sockets Layer 1 (SSL1)	1.0-2JA	[KITS.SSL1]
VSI Secure Sockets Layer (SSL)	1.4-502A	[KITS.SSL]
VSI TCP/IP	5.7-13ECO05F	[KITS.TCPIP]
VSI TCP/IP NFS Patch	5.7-ECO05A	[KITS.TCPIP]
VSI TCP/IP SSH Patch	5.7-ECO05D	[KITS.TCPIP]
VSI TDC_RT	2.3-1220	[KITS.TDC_RT]
VSI VT	2.1-L	[KITS.DECNET_PLUS]



## Appendix A

# License Management Configuration Considerations for Mixed-Producer OpenVMS Clusters

### A.1 Introduction

With the release by VSI of OpenVMS V8.4-1H1 for Integrity servers and OpenVMS V8.4-2L1 for Alpha servers, an OpenVMS cluster may contain OpenVMS operating system versions supplied by two distinct suppliers (or *producers*, in the terminology of the OpenVMS License Management Facility). Those suppliers are Hewlett Packard Enterprise (HPE) and VSI. This document describes how to configure the OpenVMS License Management Facility in a mixed-supplier cluster.

### A.2 Background

The OpenVMS License Management Facility (LMF) enforces the HPE OpenVMS and 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).

### A.3 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 Section 2.1 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.”

## A.4 Alpha LMF License Types

Alpha license PAKs come in three possible varieties although VSI uses only the Type A and Type H PAKs. To determine the required units for each category, enter the SHOW LICENSE/CHARGE command at the DCL prompt as shown below:

```
$ SHOW LICENSE/CHARGE
VMS/LMF Charge Information for node DEMO
This is a hp AlphaServer ES47 7/1150, hardware model type 2030
Type: A, Units Required: 125 (VAX/VMS Capacity or OpenVMS Unlimited or Base)
Type: B, * Not Permitted * (VAX/VMS F&A Server)
Type: C, * Not Permitted * (VAX/VMS Concurrent User)
Type: D, * Not Permitted * (VAX/VMS Workstation)
Type: E, * Not Permitted * (VAX/VMS System Integrated Products)
Type: F, * Not Permitted * (VAX Layered Products)
Type: G, * Not Permitted * (Reserved)
Type: H, Units Required: 1050 (Alpha Layered Products)
Type: I, Units Required: 1050 (Layered Products)
```

VSI has consolidated all Alpha products into two PAKS:

- The Type A licenses are covered by the ALPHA-SYSTEM PAK.
- The Type H licenses are covered by the ALPHA-LP PAK.

These two PAKs represent license groups. Use the F\$LICENSE lexical function to verify individual PAKs, as shown here:

```
$ WRITE SYS$OUTPUT F$LICENSE("OPENVMS-ALPHA", "VSI")
TRUE
$ WRITE SYS$OUTPUT F$LICENSE("UCX", "VSI")
TRUE
```

## A.5 Integrity LMF License Types

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

### Operating Environment PAKs

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 OE 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		
	VMSCUSTER		
	VMSCUSTER-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.

## A.6 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 members and VSI OpenVMS 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.

## A.7 VSI LMF Implementation (OpenVMS V8.4-1H1 onward)

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 an 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:

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 V8.4-2 or Later

Beginning with VSI OpenVMS V8.4-2 on Integrity servers and with all VSI OpenVMS versions for Alpha, LMF no longer considers VSI a synonym matching the behavior of HP OpenVMS v8.4 and earlier. VSI OpenVMS V8.4-1H1 requires a patch kit to restore the matching behavior. (See details below.) In order to correctly assign license units to the appropriate operating system, you must modify the PAKs in your LMF database to load only on operating system versions supplied by the producer in the PAK's 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, 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

The LICENSE MODIFY command accepts wildcards for the product names. For example:

```
$ LICENSE MODIFY */PRODUCER=HP/INCLUDE=(node-list...)
```

**Note:** If you are running HP products on VSI OpenVMS (See "Concerns and Questions" below), do not exclude the HP PAKs for those products from the VSI nodes. Once the VSI system is upgraded to the VSI-branded product, the licenses must be modified to load only on the appropriate systems.

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, VMSCLUSTER, VOLSHAD
IA21	VMS V8.4-1H1	MEMBER	2 cores BOE, VMSCLUSTER
IA18	VMS V8.4	MEMBER	2 cores FOE, VMSCLUSTER, 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.

### **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.

## **A.8 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](mailto: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](mailto: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.

---

Copyright © 2017 VMS Software, Inc., Bolton Massachusetts, USA

Confidential computer software. Valid license from VSI 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 VSI 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. VSI shall not be liable for technical or editorial errors or omissions contained herein.

HPE, HPE Integrity, and HPE Alpha are trademarks or registered trademarks of Hewlett Packard Enterprises.

Java is a registered trademark of Oracle and/or its affiliates.

Kerberos is a trademark of the Massachusetts Institute of Technology.

Motif is a registered trademark of The Open Group.