vms Software

Software Product Description and QuickSpecs

PRODUCT NAME: VSI OpenVMS Alpha Version 8.4-2L1   DO-DVASPQ-001

This SPD and QuickSpecs describes the VSI OpenVMS Alpha Operating System software for AlphaServer computer families.

DESCRIPTION

OpenVMS is a general purpose, multiuser operating system that runs in both production and development environments. VSI OpenVMS Alpha Version 8.4-2L1 is the latest release of the OpenVMS Alpha computing environment by VMS Software, Inc (VSI).

OpenVMS Alpha supports HPE's AlphaServer series computers. OpenVMS software supports industry standards, facilitating application portability and interoperability. OpenVMS provides symmetric multiprocessing (SMP) support for multiprocessing systems.

The OpenVMS operating system can be tuned to perform well in a wide variety of environments. This includes combinations of compute-intensive, I/O-intensive, client/server, real-time, and other environments. Actual system performance depends on the type of computer, available physical memory, and the number and type of active disk and tape drives.

The OpenVMS operating system has well-integrated networking, distributed computing, client/server, windowing, multi-processing, and authentication capabilities. It contains extensive features that promote ease-of-use, improve the productivity of programmers, facilitate system management, and provide effective operating system security.

For information about the OpenVMS Version 8.4-2L1 new features, see the *VSI OpenVMS Version 8.4-2L1 Cover Letter and Release Notes* document at: http://www.vmssoftware.com.

USER ENVIRONMENT

Users can access the OpenVMS software by using the English-like DIGITAL Command Language (DCL), the command language for OpenVMS that is supplied with the system. DCL commands provide information about the system and initiate system utilities and user programs. DCL commands take the form of a command name followed by parameters and qualifiers.

Users can enter DCL commands at a terminal or include them in command procedures. These command procedures can be run interactively or submitted to a batch queue for later processing. Information about DCL and OpenVMS utilities is available on line through the OpenVMS Help system.

For users who are familiar with the UNIX shell and utilities, an open source port of GNV is available. GNV implements a UNIX environment on OpenVMS and includes an Implementation of the UNIX shell BASH (Bourne Again Shell) and many UNIX-shell utilities.

The following tools and utilities are integrated into the OpenVMS operating system.

Text Processing

The Extensible Versatile Editor (EVE) is the default editor for OpenVMS. EVE allows users to insert, change, and delete text quickly. EVE is a full-screen editor that allows users to scroll through text on a terminal screen. EVE provides an EDT-style keypad, allowing EDT users to move easily to EVE.

Mail Utility

The Mail utility allows users to send messages to any other user on the system. Multinode operation is available if a DECnet or TCP/IP product is installed and licensed on each participating node on the network.

Command-Level Programming

Command-level programming allows users to create special files, called command procedures that contain a series of DCL commands. When users execute a command procedure, the system processes the commands in the command procedure consecutively.

User Environment Tailoring

Users can customize the computing environment with login command procedures, shorthand commands, binding of commands to function keys, command recall and editing, and process logical names.

PROGRAM DEVELOPMENT ENVIRONMENT

OpenVMS includes a comprehensive set of tools for developing programs, including: run-time libraries (RTLs), system service routines, a linker, a librarian, and a symbolic debugger.

The following tools are available to the OpenVMS programmer.

Language and Run-Time Library Support

OpenVMS includes several RTLs that provide:

- String manipulation
- Parallel processing support
- I/O routines
- I/O conversion
- Terminal-independent screen handling
- Date and time formatting routines
- Highly accurate mathematical functions
- Signaling and condition handling
- Other general-purpose functions

With OpenVMS, these routines can be called from programs written in such languages as MACRO-32, MACRO-64, Ada, BASIC, C, C++, COBOL, FORTRAN, and Pascal.

OpenVMS includes language-support libraries for extending language functionality on OpenVMS. While each language is different, all provide support for sequential file I/O, and most support direct and indexed file I/O. Language RTLs also provide support for I/O formatting, error handling, and in Fortran, the ability to read unformatted files that contain data from other vendors. RTLs are provided to support translated images created from user-mode images built on OpenVMS Alpha Version 6.1 through Version 7.3-2.

Calling Standard

All OpenVMS programming languages adhere to the common calling standard. This means that routines written in any of these languages can directly call routines written in any other language. Development of applications using multiple languages is simple and straightforward.

All user-accessible routines in the RTLs follow the appropriate platform calling standard and condition-handling conventions, and most are contained within shareable images.

System services provide access to basic operating system functions and interprocess communications as well as various control resources. Programs can call system services routines directly for security, event flag, asynchronous system trap, logical name, record and file I/O, process control, timer, time conversion, condition handling, lock management, and memory management. System services use the appropriate platform calling standard and condition-handling conventions. OpenVMS supports the execution of user-mode images created on earlier versions of OpenVMS. Typically, recompiling and relinking are not required.

MACRO Compiler

With minor modifications, VAX MACRO-32 sources can be compiled for execution on AlphaServers.

POSIX Threads Library

OpenVMS includes a user-mode, multithreading capability called POSIX Threads Library, which provides a POSIX 1003.1-1996 standard style threads interface. Additionally, the library provides an interface that is the OpenVMS implementation of Distributed Computing Environment (DCE) threads as defined by The Open Group.

POSIX Threads Library consists of run-time routines that allow the user to create multiple threads of execution within a single address space. With POSIX Threads Library Kernel Threads features enabled, POSIX Threads Library provides for concurrent processing across all CPUs by allowing a multithreaded application to have a thread executing on every CPU (on both symmetric and asymmetric multiprocessor systems). Multithreading allows computation activity to overlap I/O activity. Synchronization elements, such as mutexes and condition variables, are provided to help ensure that shared resources are accessed correctly. For scheduling and prioritizing threads, POSIX Threads Library provides multiple scheduling policies. For debugging multithreaded applications, POSIX Threads Library is supported by the OpenVMS Debugger. POSIX Threads Library also provides Thread Independent Services (TIS), which assist in the development of thread-safe APIs.

Librarian Utility

The Librarian utility permits storage of object modules, image files, macros, help files, text files, or any general record-oriented information in central, easily accessible files. Object module and image file libraries are searched by the linker when the linker finds a reference it cannot resolve in one of its input files. Macro libraries are searched by MACRO-32 and MACRO-64 when either finds a macro name that is not defined in the input file.

Hypersort

Hypersort is a portable library of user-callable routines that provide a high performance sorting capability for AlphaServer systems.

Traceback Facility

The Traceback facility is a debugging tool that provides symbolic information about call stack PCs. When an application is compiled and linked with traceback information, the Traceback facility translates stack frame addresses into routine names and line numbers and displays a symbolic traceback whenever a runtime error occurs in that application.

Debugger

The OpenVMS Debugger allows users to trace program execution, as well as display and modify register contents using the same symbols that are present in the source code. The debugger contains a heap analyzer feature that displays a graphic view of memory allocations and deallocations in real time.

System Code Debugger

The OpenVMS System Code Debugger is a kernel code debugger. It allows a system code developer to trace the execution of nonpageable system code at any interrupt priority level (IPL). Based on the OpenVMS Debugger, the System Code Debugger uses the same interface and most of the same command set.

System Dump Analyzer (SDA) Utility

In the event of a system failure, OpenVMS writes the contents of memory to a preallocated dump file. This dump file can later be analyzed using System Dump Analyzer (SDA). System dumps can either be full memory dumps, where all memory is written, or selective memory dumps, where only portions of memory in use at the time of the system failure are written. The dump file can be located on any locally connected disk. On Alpha servers, dump compression allows both full and selective dumps to be written to smaller files than required for uncompressed

dumps. Full memory dumps, if not compressed, require a dump file big enough to hold all memory. Selective memory dumps write as much of the memory in use at the time of the system failure that will fit into the dump file.

Spinlock Tracing Utility

The Spinlock Tracing Utility provides a mechanism to characterize spinlock usage and collect performance data for a given spinlock on a per-CPU basis. It can identify which spinlock is heavily used and what process is acquiring and releasing spinlocks.

Process Dumps

When an application fails, a copy of its registers and memory can be written to a data file, which can be examined using the ANALYZE PROCESS utility. This utility uses the same interface and commands as the OpenVMS Debugger to allow registers and memory to be examined. Another process can initiate the writing of the memory dump.

RMS File Utilities

Record Management Services (RMS) file utilities allow users to analyze the internal structure of an RMS file and tune the I/O, memory, space and performance parameters of the file. The RMS file utilities can also be used to create, load, and reclaim space in an RMS file. For more information about RMS, see the Operating System Environment section of this SPD.

File Differences Utility

This utility compares the contents of two files and lists those records that do not match.

Translated Image Environment (TIE) (Alpha)

OpenVMS Alpha provides an array of services that allow the operation of programs that have undergone binary translation from OpenVMS VAX images. These programs perform virtually all user-mode functions on OpenVMS Alpha and operate in combination with other programs (images) that have been translated from OpenVMS VAX or have been built using native compilers on OpenVMS Alpha. Without requiring special source code, the TIE resolves differences between the VAX and Alpha architectures, including floating- point registers, condition codes, exception handling, and ASTs. The TIE included with OpenVMS Alpha can run images that have been translated elsewhere.

NOTE: Binary translation is a means of moving complex software application executable images, where source code is not available, from one architecture and operating system to another architecture and operating system.

SYSTEM MANAGEMENT ENVIRONMENT

OpenVMS provides a set of tools and utilities that aid the system manager in configuring and maintaining an optimal system as follows:

VSI Availability Manager

VSI Availability Manager is a system management tool that enables you to monitor one or more OpenVMS nodes on an extended local area network (LAN) from either an OpenVMS Alpha system, or an OpenVMS for Integrity server system, or a PC running Windows. This tool helps system managers and analysts target a specific node or process for detailed analysis and also can resolve certain performance or resource problems. It is the multiplatform replacement for the DECamds product and includes the DECamds functionality in its capabilities.

Availability Manager has a wide-area capability whereby any system on the network supporting Availability Manager can be managed from a central console. Moreover, Availability Manager is enhanced to support Cluster over IP to manage and monitor LAN or IP path data, and IP interface for cluster communication.

The Data Collector, part of the Availability Manager product, collects system and process data on an OpenVMS node and should be installed on each node that you need to monitor (Alpha and Integrity servers).

The Data Analyzer analyzes and displays the data collected by the Data Collector, and can analyze and display data from many OpenVMS nodes simultaneously (OpenVMS Alpha nodes, and PCs running 64-bit Windows).

Hardware recommendations and related documentation are available in the *Availability Manager Installation Guide*.

Performance Data Collector

Performance data for an AlphaServer system can be gathered using the Performance Data Collector (TDC). By default, TDC periodically collects and stores data in a file that can be retrieved by user applications. A TDC Software Developers Kit (SDK) supports integration of TDC with new or existing applications and allows processing of "live" data as well as data read from files. TDC runtime software is installed with OpenVMS.

TDC runtime software (TDC_RT Version 2.3-1220) is installed with OpenVMS Version 8.4-2L1.

Additional information can be found in the *VSI TDC-RT Software Product Description* (SPD DO-VIBHAB-005).

MONITOR Utility

The Monitor utility (MONITOR) is a system management tool used to obtain information about operating system performance.

MONITOR allows you to monitor classes of system-wide performance data (such as system I/O statistics, page management statistics, and time spent in each of the processor modes) at specifiable intervals, and produce several types of output.

The MONITOR utility gets data that is sampled by OpenVMS approximately every millisecond. Since MONITOR and other measuring tools get data by sampling, the accuracy of the data may be affected for an application, which changes states at a rate close to the sampling interval. For example, you might see inaccurate CPU use reporting for an application that sleeps most of the time, but wakes up momentarily every millisecond.

Class Scheduler for CPU Scheduling

The Class Scheduler is a SYSMAN-based interface for defining and controlling scheduling classes for OpenVMS systems that allows you to designate the percentage of CPU time that a system's user may receive by placing users into scheduling classes.

Batch and Print Queuing System

OpenVMS provides an extensive batch and print capability that allows the creation of queues and the setup of spooled devices to process non-interactive workloads in parallel with timesharing or real-time jobs.

The OpenVMS batch and print operations support two types of queues: generic queues and execution queues. A generic queue is an intermediate queue that holds a job until an appropriate execution queue becomes available to initiate the job. An execution queue is a queue through which the job (either print or batch) is actually processed. Because multiple execution queues can be associated with a generic queue, OpenVMS enables load balancing across available systems in an OpenVMS Cluster system, increasing overall system throughput.

Print queues, both generic and execution, together with queue management facilities, provide versatile print capabilities, including support for various print file formats.

Accounting Utility

For accounting purposes, OpenVMS keeps records of system resource usage. Statistics include processor and memory utilization, I/O counts, print symbiont line counts, image activation counts, and process termination records. The OpenVMS Accounting utility allows you to generate reports using this data, in order to learn more about how the system is used and how it performs.

Audit Analysis Utility

For security auditing purposes, OpenVMS selectively records critical, security-relevant events in the system security audit log file. These records contain the date and time the event occurred, the identity of the associated user process, and information specific to each event type. This information helps the system manager maintain system security and deter possible intruders. The OpenVMS Audit Analysis utility allows you to generate various reports from this data.

Autoconfigure and AUTOGEN Utilities

The Autoconfigure and AUTOGEN utilities automatically configure the available devices in the system tables and set system parameters based on the peripheral and memory architecture. This eliminates the need for a traditional system generation process when the hardware configuration is expanded or otherwise modified.

The OpenVMS AUTOGEN command procedure sets several system parameters automatically by detecting the devices installed in a configuration. A feedback option allows you to generate a report of recommended parameter settings based on previous usage patterns.

Backup Utility

The Backup utility provides both full-volume and incremental file backups for file-structured, mounted volumes and volume sets. Individual files, selected directory structures, or all files on a volume set can be backed up and restored. Files can be selected by various dates (such as creation or modification) and can be backed up to magnetic tape, magnetic disk, or Write Once Read Many (WORM) optical disk. The Backup utility can also be used to restore a saveset or list the contents of a saveset.

The Backup utility has been extended to support volumes up to 2 TB. Backup utility has also been enhanced to create and restore a compressed save set. The compressed save set can be created on disks and magnetic tapes. The compression ratio depends on the data content in the files.

A Backup API is included for invoking backup routines from an executable procedure.

The Backup Manager for OpenVMS provides a screen- oriented interface to the Backup utility that assists users in performing routine backup operations. The Backup Manager is menu driven and provides:

- Access to the save, restore, and list operations without having to understand Backup command syntax
- The ability to create, modify, recall, and delete Backup Manager templates that describe the Backup save operations

Recordable CD

On supported AlphaServer systems that ship with writable CD devices (CD-RW), OpenVMS provides the capability to write once to CD-R media using an application shipping in the base operating system. For application details, please refer to the OpenVMS documentation set.

Analyze Disk Structure Utility

The Analyze Disk Structure utility compares the structure information on a disk volume with the contents of the disk, prints the structure information, and permits changes to that information. It can also be used to repair errors detected in the file structure of disks.

License Management Facility (LMF)

The License Management Facility allows the system manager to enable software licenses and to determine which software products are licensed on an OpenVMS system.

System Management Utility (SYSMAN)

The System Management utility allows system managers to define a management environment in which operations performed from the local OpenVMS system can be executed on all other OpenVMS systems in the environment or cluster. This allows multiple OpenVMS systems to be managed as easily as a single system.

SECURITY

OpenVMS provides a rich set of tools to control user access to system-controlled data structures and devices that store information. OpenVMS employs a reference monitor concept that mediates all access attempts between subjects (such as user processes) and security-relevant system objects (such as files). OpenVMS also provides a system security audit log file that records the results of all object access attempts. The audit log can also be used to capture information regarding a wide variety of other security-relevant events.

User account information, privileges and quotas associated with each user account is maintained in the system user authorization file (SYSUAF). Each user account is assigned a user name, password, and unique user identification code (UIC). To log in and gain access to the system, the user must supply a valid user name and password. The password is encoded and does not appear on terminal displays.

Users can change their password voluntarily, or the system manager can specify how frequently passwords change, along with minimum password length, `password history policy`, and the use of randomly generated passwords.

Operations

OpenVMS allows for varying levels of privilege to be assigned to different operators. Operators can use the OpenVMS Help Message utility to receive online descriptions of error messages. In addition, system-generated messages can be routed to different terminals based on their interest to the console operators, tape librarians, security administrators, and system managers.

Security auditing is provided for the selective recording of security-related events. This auditing information can be directed to security operator terminals (alarms) or to the system security audit log file (audits). Each audit record contains the date and time of the event, the identity of the associated user process, and additional information specific to each event.

OpenVMS provides security auditing for the following events:

- Login and logout
- Login failures and break-in attempts
- Object creation, access, deaccess, and deletion; selectable by use of privilege, type of access, and on individual objects
- Authorization database changes
- Network logical link connections for DECnet for OpenVMS, DECnet-Plus, DECwindows, IPC, and SYSMAN
- Use of identifiers or privileges
- Installed image additions, deletions, and replacements
- Volume mounts and dismounts
- Use of the Network Control Program (NCP) utility
- Use or failed use of individual privileges
- Use of individual process control system services
- System parameter changes
- System time changes and recalibrations

Every security-relevant system object is labeled with the UIC of its owner along with a simple protection mask. The owner UIC consists of two fields: the user field and a group field. System objects also have a protection mask that allows read, write, execute, and delete access to the object's owner, group, privileged system users, and to all other users. The system manager can protect system objects with access control lists (ACLs) that allow access to be granted or denied to a list of individual users, groups, or identifiers. ACLs can also be used to audit access attempts to critical system objects.

OpenVMS applies full protection to the following system objects:

- Common event flag clusters
- Devices
- Files
- Group global sections
- Logical name tables
- Batch/print queues
- Resource domains
- Security classes
- System global sections
- ODS-2 volumes
- ODS-5 volumes

OpenVMS provides optional security solutions to protect your information and communications:

- OpenVMS includes encryption for data confidentiality that ships as part of the operating system, thereby removing the requirement to license and install Encrypt separately. The ENCRYPT and DECRYPT commands, part of OpenVMS, support Advanced Encryption Standard (AES) file encryption with 128, 192, or 256 bit keys. AES encryption is also supported by BACKUP/ENCRYPT, allowing for the creation of encrypted tapes and savesets. The built-in encryption functionality is backward-compatible with file and backup tapes created by the former layered product Encryption for OpenVMS. This layered product featured 56-bit Data Encryption Standard (DES), which continues to function today, allowing for the decryption of archived DES encrypted data. The AES encryption functionality supports Electronic Code Book (ECB) and Cipher Block Chaining (CBC) block modes of encryption. The Cipher Feedback (CFB) and Output Feedback (OFB) 8-bit character stream modes are also supported from the command line as well as by the programmatic APIs.

- Secure Sockets Layer1 (SSL1) for OpenVMS provides secure transfer of sensitive information over the Internet

- TCP/IP allows use of an RSA host key that enables secure connectivity with newer SSH client implementations without requiring reconfiguration of the client to support the older, less-secure DSA host key types.

- Common Data Security Architecture (CDSA) is configured and initialized automatically during installation and upgrades and is required for Secure Delivery purposes and other security features. If you install a newer version of CDSA without upgrading the base operating system, you must initialize the CDSA software, using the following command. Enter the command from an account that has both SYSPRV and CMKRNL privileges (for example, the SYSTEM account). $ @SYS$STARTUP:CDSA$UPGRADE

- Kerberos for OpenVMS

- Per-thread security profiles

- External authentication

- Global and local mapping of Lightweight Directory Access Protocol (LDAP) users

- VSI code signing for OpenVMS: OpenVMS kits will be signed using VSI Code Signing Service (CSS)

NOTE: Users who are externally authenticated by their LAN Manager need only remember a single user name/password combination to gain access to their OpenVMS and LAN Manager accounts.

NOTE: Because no system can provide complete security, VSI cannot guarantee complete system security. However, VSI continues to enhance the security capabilities of its products. Customers are strongly advised to follow all industry-recognized security practices. OpenVMS recommended procedures are included in the *OpenVMS Guide to System Security*.

OPERATING SYSTEM ENVIRONMENT

Processes and Scheduling

Executable images consist of system programs and user programs that have been compiled and linked. These images run in the context of a process on OpenVMS systems. Sixty-four process priorities are recognized on OpenVMS. Priorities 0 to 15 are for time-sharing processes and applications (four is the typical default for timesharing processes). Priorities 16 to 63 are for real-time processes. Real-time processes can be assigned higher priorities to ensure that they receive processor time whenever they are ready to execute.

OpenVMS uses paging and swapping to provide sufficient virtual memory for concurrently executing processes. Paging and swapping is also provided for processes whose memory requirements exceed available physical memory.

64-Bit Virtual Addressing

The OpenVMS operating system provides support for 64-bit virtual memory addressing. This capability makes the 8 TB virtual address space available to the OpenVMS Alpha operating system and to application programs. OpenVMS applications can take advantage of 64-bit processing by using 64-bit data types supported by the compilers. For further details, see the SPDs for the OpenVMS compilers.

Very Large Memory (VLM) Features

OpenVMS provides the following additional memory management VLM features beyond those provided by 64-bit virtual addressing. These features can be used by database servers to keep large amounts of data in memory, which may result in dramatically increased runtime performance.

OpenVMS provides these VLM features:

- Memory-resident global sections

- Fast I/O for global sections

- Shared page tables

- Expandable global page table

- Reserved memory registry

DECdtm Services

DECdtm services embedded in the OpenVMS operating system support fully distributed databases using a two-phase commit protocol. The DECdtm services provide the technology and features for distributed processing, ensuring both transaction and database integrity across multiple resource managers. Updates to distributed databases occur as a single all-or-nothing unit of work, regardless of where the data physically resides. This ensures the consistency of distributed data.

DECdtm services allow applications to define global transactions that can include calls to any number of VSI data management products. Regardless of the mix of data management products used, the global transaction either commits or aborts. OpenVMS is unique in providing transaction-processing functionality with base operating system services.

DECdtm features include:

- Embedded OpenVMS system services that support the DECtp architecture, providing the features and technology for distributed transaction processing.

- Ability for multiple disjoint resources to be updated automatically. These resources can be either physically disjointed on different clusters at separate sites, or logically disjointed in different databases on the same node.

- Ability to use the X/Open Distributed Transaction Processing XA interface that enables the DECdtm transaction manager to coordinate XA-compliant resource managers (the VSI DECdtm XA Veneer), and XA-compliant transaction processing systems to coordinate DECdtm-compliant resource managers (the DECdtm XA Gateway).

- Robust application development. Applications can be written to ensure that data is never in an inconsistent state, even in the event of system failures.

- Ability to be called using any VSI TP monitor or database product. This is useful for applications using several VSI database products.

Interprocess Communication

OpenVMS provides the following facilities for applications that consist of multiple cooperating processes:

- Mailboxes as virtual devices that allow processes to communicate with queued messages.

- Shared memory sections on a single processor or an SMP system that permit multiple processes to access shared address space concurrently.

- Common event flags that provide simple synchronization.

- A lock manager that provides a more comprehensive enqueue/dequeue facility with multilevel locks, values, and asynchronous system traps (ASTs).

- Intracluster communication services through which two processes running on the same system or on different OpenVMS Cluster nodes can establish a connection and exchange data.

- Logical names through which one process can pass information to other processes running on the same system or on different OpenVMS Cluster nodes.

- Network interprocess communication is available via TCP/IP Services and DECnet-Plus.

Symmetric Multiprocessing (SMP)

OpenVMS provides symmetric multiprocessing (SMP) support for AlphaServer multiprocessor systems. SMP is a form of tightly coupled multiprocessing in which all processors perform operations simultaneously. All processors perform operations in all OpenVMS access modes, user, supervisor, executive, and kernel.

OpenVMS SMP configurations consist of multiple CPUs executing code from a single shared memory address space. Users and processes share a single copy of OpenVMS address space. SMP also provides simultaneous shared access to common data in global sections to all processors. OpenVMS SMP selects the CPU where a process will run based on its priority and in special cases as directed by the application. OpenVMS uses a specialized scheduling algorithm when running a non-uniform memory access (NUMA) platform.

SMP support is an integral part of OpenVMS and is provided to the user transparently. Because an SMP system is a single system entity, it is configured into a network and OpenVMS Cluster configurations as a single node.

The maximum number of supported CPUs in an SMP configuration is 32.

Networking Facilities

OpenVMS provides device drivers for local area network (LAN) adapters. Application programmers can use the QIO system service to communicate with other systems connected via the LAN in accordance with the IEEE Standard for Ethernet, IEEE Std 802.3™-2015.

OpenVMS supports the following networking products:

- VSI TCP/IP Services for OpenVMS, the industry-standard set of protocols for interoperating between different operating systems

- VSI DECnet-Plus

- VSI DECnet Phase IV

These networking products are described in this SPD under Associated Products.

Terminal Server Products

Terminal server products provide network access to OpenVMS from terminal (serial) based devices. When used in an OpenVMS Cluster environment, terminal servers distribute users across the available server systems at login time. OpenVMS can also establish a connection to other devices such as printers or other serially attached devices attached to such terminal servers.

Universal Serial Bus Support

OpenVMS supports the Universal Serial Bus (USB) technology. Support for the USB interconnect enables OpenVMS systems to connect to multiple supported USB devices using a single USB cable. OpenVMS supports one USB keyboard and mouse on systems that are supported by OpenVMS and have USB hardware and a graphics controller.

OpenVMS provides a USB configuration tool called UCM that can be used to track USB configuration changes like plug and unplug events. UCM can also be used to restrict the automatic addition of specific devices and classes of devices. The UCM event log is used by VSI to help diagnose problems with USB devices.

OpenVMS Alpha USB support is limited to low and full speed devices only, and currently supports only the AlphaServer ES47.

Reliability

OpenVMS handles hardware errors as transparently as possible while maintaining data integrity and providing sufficient information to diagnose errors. The system limits the effects of an error by first determining if the error is fatal. If the error occurs in system context, the current OpenVMS system shuts down. If the error is not fatal, the system recovers actions pertinent to the error and continues the current operation.

In all cases, information relevant to the error is written to the error log file for later analysis. Hardware errors include the following categories:

- **Processor errors.** These include processor soft errors, processor hard errors, processor machine checks, and adapter errors.

- **Memory errors.** These can be unrecoverable (hard) errors or recoverable (soft) errors. The system examines memory at startup time and does not use any bad pages. During system operation, the system corrects all single-bit memory errors for those systems with error correction code (ECC) memory.

- **Correctible memory errors.** A primary cause of these correctible memory errors is alpha particle radiation. On some processors, when correctible memory errors occur, the memory controller corrects only the data returned to the CPU or I/O controller. The actual data in memory is left with the error intact. Subsequent read operations cause correction cycles to occur and, in most cases, an interrupt to report the error. On many of these processors, OpenVMS monitors the occurrence of correctible memory errors and, in almost all cases, is able to remove the error condition by rewriting the data in memory. Rewriting the data causes the data to be corrected in that memory location.

Other failures include:

- Operating system errors (system-detected inconsistencies or architectural errors in system context)

- User errors

- I/O errors

The system logs all processor errors, all operating system errors detected through internal consistency checks, all double-bit memory errors (and a summary of corrected single-bit memory errors), and most I/O errors.

If the system is shut down because of an unrecoverable hardware or software error, a dump of physical memory is written. The dump includes the contents of the processor registers. The OpenVMS System Dump Analyzer (SDA) utility is provided for analysis of memory dumps.

OpenVMS supports CPU Component Indictment, also called Dynamic Processor Resilience (DPR). When certain error conditions persist a CPU will be stopped and no longer used by the running system. Use of this feature is controlled by the System Manager via SYS$MANAGER:SYS$INDICTMENT_POLICY.COM.

Input/Output

The QIO system service and other related I/O services provide a direct interface to the operating system's I/O routines. Services are available from within most OpenVMS programming languages and can be used to perform low-level I/O operations efficiently with a minimal amount of system overhead for time-critical applications.

Device drivers execute I/O instructions to transfer data to and from a device and to communicate directly with an I/O device. Each type of I/O device requires its own driver. VSI supplies drivers for all devices supported by the OpenVMS operating system and provides QIO system service routines to access the special features available in many of these devices.

OpenVMS supports a variety of disk and tape peripheral devices, as well as terminals, networks, and mailboxes (virtual devices for interprocess communication), and more general I/O devices.

I/O Performance Features

Fast I/O provides a suite of additional system services that applications can use to improve I/O throughput. Fast I/O services minimize the CPU resources required to perform I/O.

Fast Path provides a streamlined mainline code path through the I/O subsystem to improve both uniprocessor and multiprocessor I/O performance. On multiprocessor systems, Fast Path allows all CPU processing for specific I/O adapters to be handled by a specific CPU. This can significantly lower the demands on the primary CPU and increase the I/O throughput on multiprocessor systems with multiple I/O ports. No user application changes are needed to take advantage of Fast Path. Fast Path can be utilized by the $QIO system service or the Fast I/O services.

Extended File Cache (XFC)

The Extended File Cache (XFC) is a virtual block data cache provided with OpenVMS. Similar to the Virtual I/O Cache, the XFC is a clusterwide, file system data cache. Both file system data caches are compatible and coexist in the OpenVMS Cluster.

The XFC improves I/O performance with the following features that are not available with the virtual I/O cache:

* Read-ahead caching

* Automatic resizing of the cache

* Larger maximum cache size

* No limit on the number of closed files that can be cached

* Control over the maximum size of I/O that can be cached

* Control over whether cache memory is static or dynamic

XFC caching attributes of volume can be dynamically modified eliminating the need to dismount the volume.

Virtual I/O Cache

OpenVMS Alpha provides a standalone or clusterwide, file-oriented disk cache. Applications benefit from the advantages of the virtual I/O cache without any special coding. The virtual I/O file-caching algorithm is chosen based on the type of clusterwide access currently in progress. Virtual I/O caching reduces current and potential I/O bottlenecks within OpenVMS systems. It reduces the number of I/Os to the disk subsystem, thereby reducing systemwide bottlenecks.

Record Management Services (RMS)

RMS is a set of I/O services that helps application programs to process and manage files and records. Although intended to provide a comprehensive software interface to mass storage devices, RMS also supports device-independent access to unit-record devices.

RMS supports sequential, relative, and indexed file organizations in fixed-length or variable-length record formats. RMS also supports byte stream formats for sequential file organization.

RMS record access modes provide access to records in four ways:

* Sequentially

* Directly by key value

* Directly by relative record number

* Directly by record file address

RMS also supports block I/O operations for various performance-critical applications that require user-defined file organizations and record formats.

RMS promotes safe and efficient file sharing by providing multiple file access modes and automatic record locking (where applicable). RMS offers the options of enabling global buffers for buffer sharing by multiple processes.

RMS utilities aid file creation and record maintenance. These utilities convert files from one organization and format to another; restructure indexed files for storage and access efficiency; and reclaim data structures within indexed files. These utilities also generate appropriate reports.

For systems that have DECnet or DECnet-Plus installed, RMS provides a subset of file and record management services to remote network nodes. Remote file operations are generally transparent to user programs.

Commands such as EDIT, CREATE, COPY, TYPE, and PRINT allow users to manipulate RMS records within RMS files at the DCL command level.

Disk and Tape Volumes

The system manager can organize disk volumes into volume sets. Volume sets can contain a mix of disk device types and can be extended by adding volumes. Within a volume set, files of any organization type can span multiple volumes. Files can be allocated to the set as a whole (the default) or to specific volumes within the set. Optionally, the system manager can allocate portions of indexed files to specific areas of a single disk or to specific volumes in a volume set.

The system manager can place quotas on a disk to control the amount of space individual users can allocate.

Quota assignment is made by UIC and can be controlled for each individual volume set in the system (or for each individual volume if the volume is not part of a set).

The system manager can cache disk structure information in memory to reduce the I/O overhead required for file management services. Although not required to do so, users can preallocate space and control automatic allocation. For example, a file can be extended by a given number of blocks, contiguously or noncontiguously, for optimal file system performance.

The system applies software validity checks and checksums to critical disk structure information. If a disk is improperly dismounted because of user error or system failure, the system rebuilds the disk's structure information automatically the next time the disk is mounted. The system detects bad blocks and prevents their reuse once the files to which the blocks were allocated are deleted. On DIGITAL Storage Architecture (DSA) disks, the disk controller detects and replaces bad blocks automatically.

The system provides 255 levels of named directories and subdirectories whose contents are alphabetically ordered. Device and file specifications follow VSI conventions. Users can use logical names to abbreviate the specifications and to make application programs device and file name independent. Users can assign a logical name to an entire specification, to a portion of a specification, or to another logical name.

OpenVMS supports multivolume magnetic tape files with transparent volume switching. Access positioning is done either by file name or by relative file position.

ASSOCIATED PRODUCTS

The products in this section are licensed with, and included as, part of the OpenVMS Alpha Operating System.

VSI Galaxy Software Architecture on OpenVMS Alpha

VSI Galaxy Software Architecture on OpenVMS Alpha is a System Integrated Product (SIP). By running multiple instances of OpenVMS in a single computer or hard partition, an OpenVMS Galaxy computing environment gives you quantum improvements in:

- Compatibility—Existing applications run without changes.
- Availability—Presents opportunities to upgrade software and expand system capacity without downtime.
- Scalability—Offers scaling alternatives that improve performance of SMP and cluster environments.
- Adaptability—Physical resources can be dynamically reassigned to meet changing workload demands.
- Cost of ownership—Fewer computer systems reduce system management requirements, floor space, and more.

For more information about how to create, manage, and use an OpenVMS Galaxy computing environment, see the *OpenVMS Alpha Partitioning and Galaxy Guide*.

VSI OpenVMS Cluster Software

VSI OpenVMS Cluster software for Alpha server systems provides a highly integrated OpenVMS computing environment that is distributed over multiple systems, separated in distance measured from feet up to 500 miles, containing up to 96 nodes.

OpenVMS Cluster systems and storage communicate using a combination of the following interconnects:

- Ethernet

- Small Computer Systems Interface (SCSI) (Storage Only)

- Fibre Channel (Storage Only)

VSI TCP/IP Services (minimum Version 5.7-13ECO05B) is needed if using IP for cluster communication.. For more information, see the *Guidelines for OpenVMS Cluster Configurations*.

Applications running on one or more nodes in an OpenVMS Cluster system share resources in a coordinated manner. While updating data, the OpenVMS Cluster software synchronizes access to shared resources, preventing multiple processes on any node in the cluster from uncoordinated access to shared data. This coordination ensures data integrity during concurrent update transactions.

VSI supports mixed-architecture and mixed-version clusters that contain both Alpha and Integrity server systems.

Cross-architecture booting (booting an Integrity satellite node from an Alpha boot server and vice-versa) is not supported.

For more information, see the *VSI OpenVMS Cluster Software Product Description* (SPD DO-VIBHA*-033).

VSI Volume Shadowing for OpenVMS

VSI Volume Shadowing performs disk-mirroring operations using a redundant array of independent disks (RAID-1) storage strategy.

Volume Shadowing for OpenVMS provides high data availability for disk devices by ensuring against data loss that results from media deterioration or controller or device failure. This prevents storage subsystem component failures from interrupting system or application tasks. It also allows users to dynamically add new storage to an existing environment.

For more information, see the *VSI Volume Shadowing for OpenVMS Software Product Description* (SPD DO-VIBHAA-031).

VSI RMS Journaling for OpenVMS

VSI RMS Journaling, a System Integrated Product (SIP), enables a system manager, user, or application to maintain the data integrity of RMS files in the event of a number of failure scenarios. These journaling products protect RMS file data from becoming lost or inconsistent.

RMS Journaling provides the following three types of journaling:

- **After-image journaling.** Allows users to reapply modifications that have been made to a file. This type of journaling allows users to recover files that are inadvertently deleted, lost, or corrupted.

- **Before-image journaling.** Allows users to reverse modifications that have been made to a file. This type of journaling allows users to return a file to a previously known state.

- **Recovery-unit journaling.** Allows users to maintain transaction integrity. A transaction can be defined as a series of file updates on one or more files. If any failure occurs during the transaction, recovery-unit journaling rolls back the partially completed transaction to its starting point.

For more information, see the *RMS Journaling for OpenVMS Software Product Description* (SPD DO-VIBHAA-030).

VSI TCP/IP Services for OpenVMS

VSI TCP/IP Services for OpenVMS, a System Integrated Product (SIP), is VSI's industry-standard implementation of the TCP/IP and NFS networking protocols on the OpenVMS platform. TCP/IP Services for OpenVMS provides interoperability and resource sharing among systems running OpenVMS, UNIX, Windows, and other operating systems that support TCP/IP. TCP/IP provides a comprehensive suite of functions and applications that support industry-standard protocols for heterogeneous network communications and resource sharing. TCP/IP Services for OpenVMS provides a full TCP/IP protocol suite including IP/multicasting, dynamic load balancing, rlogin proxy, network file access, remote terminal access, remote command execution, remote printing, mail, application development, Post Office Protocol (POP), SNMP Extensible agent (eSNMP), and Finger Utility. TCP/IP Version 5.7 also enables packet processing Engine (PPE), FTP anonymous light and stream control transmission protocol (SCTP) for its customers. TCP/IP allows use of an RSA host key that enables secure connectivity with newer SSH client implementations without requiring reconfiguration of the client to support the older, less-secure DSA host key types.

For further information, see the *TCP/IP Services for OpenVMS Software Product Description* (SPD DO-VIBHAA-029).

VSI DECnet-Plus and VSI DECnet Software

VSI DECnet-Plus and VSI DECnet for OpenVMS are System Integrated Products (SIPs). Note that only one version of DECnet can be active on a single system at one time. Both DECnet and DECnet-Plus allow OpenVMS systems to participate in network task-to-task communications for the purposes of file transfer, copy, and printing, application execution, etc.

DECnet-Plus offers task-to-task communications, file management, downline system and task loading, network command terminals, and network resource sharing capabilities as defined in the DIGITAL Network Architecture (DNA) Phase V protocols. DECnet-Plus provides the newest DECnet features such as extended addressing and downline load performance enhancements. DECnet-Plus integrates DECnet and OSI protocols and provides a linkage to TCP/IP using Request for Comments (RFC) 1006 and RFC 1859. DECnet and OSI applications can be run over DECnet (NSP), OSI (CLNS), and TCP/IP transports.

For further information, see the *VSI DECnet-Plus for OpenVMS Software Product Description* (SPD DO-VIBHAA-023), or the *VSI DECnet for OpenVMS Software Product Description* (SPD DO-VIBHAA-024).

VSI DECram for OpenVMS

VSI DECram for OpenVMS is a disk device driver that improves I/O performance by allowing an OpenVMS system manager to create pseudo disks (RAMdisks) that reside in main memory. Frequently accessed data can be accessed much faster from a DECram device than from a physical disk device. RAMdisks can be accessed through the file system just as physical disks are accessed, requiring no change to application or system software.

Because main memory is allocated for the DECram device, extra memory is generally required. The OpenVMS system manager can designate the amount of memory dedicated to the DECram devices and the files that will be stored on it.

For more information, see the *VSI DECram for OpenVMS Software Product Description* (SPD DO-VIBHAB-005).

VSI DECwindows Motif for OpenVMS

VSI DECwindows, a System Integrated Product (SIP), provides support for OSF/Motif, a standards-based graphical user interface, and the X user interface (XUI) in a single, runtime and development environment. DECwindows Motif displays the OSF/Motif user interface. Because both Motif and XUI are based on X.org X Window System, applications written with either toolkit will run regardless of which environment the user selects.

For more information, see the *VSI DECwindows Motif for OpenVMS Software Product Description* (SPD DO-VIBHAA-026).

CONFORMANCE TO STANDARDS

OpenVMS is based on the following public, national, and international standards.

Distributed Computing Environment (DCE) Support

DCE for OpenVMS provides a set of the distributed computing features specified by The Open Group's DCE, as well as tools for application developers. With DCE, The Open Group has established a standard set of services and interfaces that facilitate the creation, use, and maintenance of client/server applications. DCE for OpenVMS serves as the basis for an open computing environment where networks of multivendor systems appear as a single system to the user. Because DCE makes the underlying networks and operating systems transparent, application developers can easily build portable, interoperable client/server applications. Users can locate and share information safely and easily across the entire enterprise. DCE for OpenVMS supplies system managers with a set of tools to consistently manage the entire distributed computing environment, while assuring the integrity of the enterprise.

DCE for OpenVMS currently consists of the following products:

- DCE Run-Time Services for OpenVMS

- DCE Application Developers' Kit for OpenVMS

- DCE Cell Directory Service (CDS)
- DCE Security Server

For more details, see the *VSI Distributed Computing Environment (DCE) for OpenVMS Software Product Description* (SPD DO-VIBHAA-027).

Support for OSF/Motif and X Window System Standards

DECwindows Motif provides support for OSF/Motif, a standards-based graphical user interface. DECwindows Motif also provides support for the X Consortium's X Window System, Version 11, Release 6 (X11R6) server and the Version 11, Release 5 (X11R5) client.

Standards Supported by OpenVMS

The OpenVMS operating system is based on the following public, national, and international standards. These standards are developed by the American National Standards Institute (ANSI), U.S. Federal Government (responsible for FIPS), Institute of Electrical and Electronics Engineers (IEEE), and the International Organization for Standardization (ISO). The following information may be useful in determining responsiveness to stated conformance requirements as enabled in particular commercial and/or government procurement solicitation documents.

• ANSI X3.4-1986: American Standard Code for Information Interchange

• ANSI X3.22-1973: Recorded Magnetic Tape (800 BPI, NRZI)

• ANSI X3.27-1987: File Structure and Labeling of Magnetic Tapes for Information Interchange

• ANSI X3.298: Limited support. Information Technology—AT Attachment-3 Interface (ATA-3)

• ANSI X3.39-1986: Recorded Magnetic Tape (1600 BPI, PE)

• ANSI X3.40-1983: Unrecorded Magnetic Tape

• ANSI X3.41-1974: Code Extension Techniques for Use with 7-bit ASCII

• ANSI X3.42-1975: Representation of Numeric Values in Character Strings

• ANSI X3.54-1986: Recorded Magnetic Tape (6250 BPI, GCR)

• ANSI X3.131-1986 (SCSI I): Small Computer System Interface

• ANSI X3.131-1994 (SCSI II): Small Computer System Interface

• ANSI/IEEE 802.2-1985: Logical Link Control

• ANSI/IEEE 802.3-1985: Carrier Sense Multiple Access with Collision Detection

• FIPS 1-2: Code for Information Interchange, Its Representations, Subsets, and Extensions
  NOTE: 1-2 includes ANSI X3.4-1977(86)/FIPS 15; ANSI X3.32-1973/FIPS 36; ANSI X3.41-1974/FIPS 35; and FIPS 7.

• FIPS 3-1/ANSI X3.22-1973: Recorded Magnetic Tape Information Interchange (800 CPI, NRZI)

• FIPS 16-1/ANSI X3.15-1976: Bit Sequencing of the Code for Information Interchange in Serial-by-Bit Data Transmission

  NOTE: FED STD 1010 adopts FIPS 16-1.

• FIPS 22-1/ANSI X3.1-1976: Synchronous Signaling Rates Between Data Terminal and Data Communication Equipment

  NOTE: FED STD 1013 adopts FIPS 22-1.

• FIPS 25/ANSI X3.39-1986: Recorded Magnetic Tape for Information Interchange (1600 CPI, Phase Encoded)

• FIPS 37/ANSI X3.36-1975: Synchronous High- Speed Data Signaling Rates Between Data Terminal Equipment and Data Communication Equipment

  NOTE: FED STD 1001 adopts FIPS 37.

• FIPS 50/ANSI X3.54-1986: Recorded Magnetic Tape for Information Interchange, 6250 CPI (246 CPMM), Group Coded Recording

• FIPS 79/ANSI X3.27-1987: Magnetic Tape Labels and File Structure for Information Interchange

• FIPS 86/ANSI X3.64-1979: Additional Controls for Use with American National Standard Code for Information Interchange

  NOTE: Other FIPS are not applicable.

  NOTE: Information regarding interchangeability of ANSI and FED standards with FIPS is contained in ''ADP Telecommunications Standards Index,'' July 1988, published and maintained by the General Services Administration.

• ISO 646: ISO 7-bit Coded Character Set for Information Exchange

• ISO 1001: File Structure and Labeling of Magnetic Tapes for Information Interchange

- ISO 1863: Information Processing — 9-track, 12, 7 mm (0.5 in) wide magnetic tape for information interchange recorded at 32 rpmm (800 rpi)

- ISO 1864: Information Processing — Unrecorded 12, 7 mm (0.5 in) wide magnetic tape for information interchange — 35 ftpmm (800 ftpi) NRZI, 126 ftpmm (3 200 ftpi) phase encoded and 356 ftmm (9 042 ftpi), NRZI

- ISO 2022: Code Extension Techniques for Use with ISO 646

- ISO 3307: Representations of Time of the Day

- ISO 3788: Information Processing — 9-track, 12, 7 mm (0.5 in) wide magnetic tape for information interchange recorded at 63 rpmm (1 600 rpt), phase encoded

- ISO 4873: 8-Bit Code for Information Interchange — Structure and Rules for Implementation

- ISO 5652: Recorded Magtape (6250)

- ISO 6429: Control Functions for Coded Character Sets

- ISO 9316: 1989 (SCSI-1) Small Computer System Interface

- ISO 9660: Information Processing — Volume and file structure of CD–ROM for information exchange

- ISO 10288: 1994 (SCSI-2) Small Computer System Interface

INSTALLATION

VSI OpenVMS Alpha is distributed as a downloadable image, available from the VMS Software, Inc, website. Procedures to set up the system disk and to prepare the system for day-to-day operations are provided in the *VSI OpenVMS Alpha V8.4-2L1 Upgrade Manual*. The procedures use the POLYCENTER Software Installation (PCSI) utility to configure and upgrade the OpenVMS operating system.

Network Upgrade

InfoServer network booting is supported for OpenVMS upgrades on hardware platforms that support OpenVMS.

OpenVMS Alpha systems can use the OpenVMS InfoServer software application or the traditional InfoServer hardware system that is independent of OpenVMS. You can use the OpenVMS InfoServer software application on all OpenVMS systems that support a DVD drive. This support provides the additional advantage of allowing a network administrator to boot multiple OpenVMS systems on the network from a single copy of the OpenVMS distribution media.

OpenVMS Alpha systems can boot from a CD on DVD drives (DVD drives support both DVDs and CDs) or from a virtual CD/DVD drive on the LAN using the OpenVMS InfoServer software application.

Any configuration procedures that might have been performed for network booting using an InfoServer hardware system (traditionally used by Alpha systems) are not valid for the OpenVMS Alpha InfoServer application.

For additional information, see the *VSI OpenVMS Alpha Version 8.4-2L1 Upgrade Manual*.

POLYCENTER Software Installation

The PCSI utility simplifies the installation and management of OpenVMS products. It is used to install, update, and uninstall software products that have been prepared with the utility. In addition, the utility provides a database to track the installation, reconfiguration, and uninstallation of software. For products installed with other installation technologies, the utility provides a mechanism for adding information about them into the product database. The utility also provides the ability to manage dependencies between products during the installation process.

For software providers, the PCSI utility simplifies the task of packaging software by providing a simple, declarative language for describing material for the installation kit and defining how it is installed. The utility handles the functions, while the developer instructs the utility what to do. This significantly reduces the complexity and time to develop installation procedures. The language allows the developer to easily specify dependencies on other software, manage objects in the execution environment (such as files and directories), and anticipate and resolve conflict before it occurs. The utility also significantly simplifies the packaging of multiple software products into one logical product suite.

Use the PCSI utility to install the operating system and to install layered products that are compliant with the POLYCENTER utility.

All of the software product kits included with OpenVMS Version 8.4-2L1 are signed using Secure Delivery. A notable exception is the OpenVMS Operating System (the OpenVMS product) because it is shipped in bootable form, not as a single file kit that is signed.Layered product kit validation is performed when you install the layered product, after the OpenVMS installation or upgrade has completed.

The PRODUCT SHOW HISTORY command displays the validation status of installed products and identifies those that were installed from unsigned kits or were installed prior to the availability of the Secure Delivery functionality.

VMSINSTAL

OpenVMS includes the VMSINSTAL facility to handle the installation of optional supplied software products that have not been converted to use the POLYCENTER Software Installation utility.

Test Package and Diagnostics

OpenVMS includes a User Environment Test Package (UETP), which verifies that the OpenVMS operating system is properly installed and ready for use on the customer's systems.

You can run diagnostics on individual devices during normal system operation. Certain critical components can operate in degraded mode.

DISK SPACE REQUIREMENTS

Operating System Disk Space Requirements

The minimum disk space required for OpenVMS Alpha is 3 GB. The disk space requirements for OpenVMS vary according to which options are installed:

File Category	Space Used
Minimum OpenVMS files	2.7 GB
DECwindows Support	89 MB
Full DECwindows Motif (optional)	109 MB
DECnet-Plus	49 MB
Other optional OpenVMS files	84 MB
Paging file (required)	38 MB
Swap file (suggested)	1 MB
Dump file (optional)	28 MB
Safeguard for upgrading	50 MB
Total	3 GB

NOTE: The minimum OpenVMS files listed in the table will allow you to run with minimal functionality. Not all OpenVMS commands and utilities will function fully as documented in this minimum configuration. Not all VSI and other layered products will work in this minimum configuration.

The minimum OpenVMS files are for a system configuration where all optional features have been declined during the initial installation. For most applications, this is not a realistic OpenVMS environment.

The paging, swap, and dump file requirements are the minimum for a system with 64 MB of main memory. Additional memory in most cases adds to the space needed for these files, as will particular needs of your application. With careful system management it is possible to use the paging file space as a temporary dump file.

For an OpenVMS Cluster system disk, paging, swap, and dump files cannot be shared between nodes, so the files must either be duplicated on the system disk or located on some other disk.

DECwindows Motif for OpenVMS Alpha Disk Space Requirements

To support full OpenVMS Alpha and full DECwindows Motif for OpenVMS Alpha, a system disk with at least 550 MB is recommended. However, a subset of the DECwindows Motif environment can be installed. The disk space required for the installation of DECwindows Motif is 159 MB. The permanent amount of space used is 145 MB. An additional 33 MB is needed to install the DECwindows X11 Display Server and associated files. (The DECwindows X11 Display Server and associated files are included in the OpenVMS Alpha operating system media.) These disk space requirements are in addition to the disk space required for the OpenVMS Alpha operating system, as indicated in the OpenVMS Alpha Disk Space Requirements table.

Installation of the DECwindows Motif layered product gives you the option of installing any or all of the following components:

- **Run-time support** (base kit) – 33 MB. This section provides support for running DECwindows Motif for OpenVMS Alpha applications on Alpha compute servers and is a required part of the installation.
- **New Desktop** — 24 MB. This is an optional component that allows use of the New Desktop environment. It includes applications and application programming interfaces (APIs).
- **DECwindows desktop** — 11 MB. This component is also optional, but you should install either the New Desktop or the DECwindows desktop to create a usable system. The DECwindows desktop is the user interface that was included in previous versions of DECwindows Motif and includes the DECwindows Session Manager, FileView, and the Motif Window Manager.
- **Programming support** — 32 MB. This section includes support for the C, C++, Fortran, and Pascal programming languages. If you install a subset of languages, the amount of disk space required will be less.
- **Example files** — approximately 26 MB.
- **Translated image support** – approximately 20 MB

Layered Product Disk Space and Memory Space Requirements

In addition to the disk space used directly by VSI or third-party layered products, there may be additional space used to store information from those products in OpenVMS help libraries, command tables, object libraries, and elsewhere. The amount of additional disk space required cannot be exactly predicted due to the possibility of recovering unused space already existing in those library files. Unusually large modules contributed by layered products can also affect the amount of space required when upgrading to a new version of the OpenVMS operating system.

MEMORY SPACE REQUIREMENTS

Operating System and Layered Product Memory Space Requirements

The minimum amount of memory required to install, boot, and log into an OpenVMS Alpha system is 64 MB. Additional memory may be required to ensure satisfactory performance for the following:

- Particular applications or number of users
- Particular hardware configurations

Refer to layered product documentation for any product-specific memory space requirements that may exist.

DISTRIBUTION METHOD

VSI OpenVMS Alpha is available via electronic download only, from the VSI website. The VSI OpenVMS Alpha electronic download images include the operating system binaries, layered product binaries, Open Source binaries, and documentation.

DOCUMENTATION

VSI OpenVMS Alpha Version 8.4-2L1 includes the following documents which are available on the VSI website:

- *VSI OpenVMS Alpha Version 8.4-2L1 Cover Letter and Release Notes*
- *VSI OpenVMS Alpha Version 8.4-2L1 Upgrade Manual*
- *VSI OpenVMS Alpha Version 8.4-2L1 Software Product Description and QuickSpecs*

GROWTH CONSIDERATIONS

The minimum hardware and software requirements for any future version of this product may be different from the requirements for the current version.

ORDERING AND LICENSING INFORMATION

OpenVMS Alpha Licensing Information

Purchase of a VSI OpenVMS Alpha support contract provides a license that enables the VSI OpenVMS Alpha operating system as well as all VSI OpenVMS Alpha layered products. The VSI OpenVMS operating system software, layered product software, and online documentation are available via download from the VSI OpenVMS website, with purchase of a support contract.

The following tables list licensing tiers available for VSI OpenVMS Alpha along with the servers included in each tier.

Table 1
VSI OpenVMS Alpha Workgroup Tiers

Workgroup Tiers		
Part Number	*Description*	*Workgroup Servers*
SP-ALPH01-D30	VSI VMS Alpha Workgroup License + 3 Year Platinum 24x7 Support	DS10 DS15 DS20 DS20E DS25
SP-ALPH01-D10	VSI VMS Alpha Workgroup License + 1 Year Platinum 24x7 Support	
SP-ALPH01-B30	VSI VMS Alpha Workgroup License + 3 Year Silver 24x7 Support	
SP-ALPH01-B10	VSI VMS Alpha Workgroup License + 1 Year Silver 24x7 Support	
SP-ALPH01-A30	VSI VMS Alpha Workgroup License + 3 Year Bronze 9x5 Support	
SP-ALPH01-A10	VSI VMS Alpha Workgroup License + 1 Year Bronze 9x5 Support	

Table 2
VSI OpenVMS Alpha Departmental Tiers

Departmental Tiers		
Part Number	*Description*	*Departmental Servers*
SP-ALPH02-D30	VSI VMS Alpha Departmental License + 3 Year Platinum 24x7 Support	ES45 ES47
SP-ALPH02-D10	VSI VMS Alpha Departmental License + 1 Year Platinum 24x7 Support	
SP-ALPH02-B30	VSI VMS Alpha Departmental License + 3 Year Silver 24x7 Support	
SP-ALPH02-B10	VSI VMS Alpha Departmental License + 1 Year Silver 24x7 Support	
SP-ALPH02-A30	VSI VMS Alpha Departmental License + 3 Year Bronze 9x5 Support	
SP-ALPH02-A10	VSI VMS Alpha Departmental License + 1 Year Bronze 9x5 Support	

Table 3
VSI OpenVMS Alpha Enterprise Tiers

Enterprise Tiers		
Part Number	*Description*	*Enterprise Servers*
SP-ALPH03-D30	VSI VMS Alpha Enterprise License + 3 Year Platinum 24x7 Support	To be announced
SP-ALPH03-D10	VSI VMS Alpha Enterprise License + 1 Year Platinum 24x7 Support	
SP-ALPH03-B30	VSI VMS Alpha Enterprise License + 3 Year Silver 24x7 Support	
SP-ALPH03-B10	VSI VMS Alpha Enterprise License + 1 Year Silver 24x7 Support	
SP-ALPH03-A30	VSI VMS Alpha Enterprise License + 3 Year Bronze 9x5 Support	
SP-ALPH03-A10	VSI VMS Alpha Enterprise License + 1 Year Bronze 9x5 Support	

Licensing Tier Support Levels

VSI OpenVMS Alpha license and support contracts are available in three levels – Platinum, Silver, and Bronze. The levels provide a choice of support length and coverage. All levels are available in 3-year or 1-year terms.

- **Platinum Level**:
  - Dedicated phone number provided on sign-up
  - Right to new minor releases of a VSI product. **NOTE**: Customers who purchase this Service Tier for a Term of 3 years or greater shall have the right to freely upgrade to any and all versions of the Software, regardless of whether such version(s) is deemed minor or major.
  - Access to all software updates/patches provided between numbered releases
  - Technical Support 24 hours a day, 7 days a week
  - Remote Technical Account Manager (TAM)
  - Periodic VSI software patch analysis
  - Periodic VSI software system health checks
  - Periodic VSI updates
  - Immediate response from initial customer contact.
- **Silver Level**:
  - Standard support phone number; email access to VSI Support.
  - Right to new minor releases of a VSI product. **NOTE**: Customers who purchase this Service Tier for a Term of 3 years or greater shall have the right to freely upgrade to any and all versions of the Software, regardless of whether such version(s) is deemed minor or major.
  - Access to all software updates/patches provided between numbered releases
  - Technical Support 24 hours a day, 7 days a week
  - Periodic VSI updates
  - Two (2) hour response from initial customer contact.
- **Bronze Level**:
  - Standard support phone number; email access to VSI Support.
  - Right to new minor releases of a VSI product. **NOTE**: Customers who purchase this Service Tier for a Term of 3 years or greater shall have the right to freely upgrade to any and all versions of the Software, regardless of whether such version(s) is deemed minor or major.
  - Access to all software updates/patches provided between numbered releases
  - Technical Support 9 hours a day, 5 days a week
  - Periodic VSI updates
  - Four (4) hour response from initial customer contact.

OpenVMS Alpha Ordering Information

Contact VSI at info@vmssoftware.com to determine which licensing tier is appropriate for your business needs. The VSI OpenVMS Alpha operating system software, layered product software, and online documentation are available via download from the VSI OpenVMS website.

OpenVMS Hardcopy Documentation

VSI is in the process of updating the OpenVMS Hardcopy documentation set; availability is TBD. The VSI OpenVMS Alpha download website contains the documents you need to get started with the OpenVMS Alpha operating system.

SOFTWARE PRODUCT SERVICES

A variety of service options are available from VSI. For more information, contact your local VSI account representative or distributor. Information is also available from:

http://www.vmssoftware.com.

SOFTWARE LICENSING

The OpenVMS operating system software is furnished under the licensing provisions of VSI's Standard Terms and Conditions. The VSI OpenVMS License provides the right to use only the VSI OpenVMS features of the current or prior versions of the VSI OpenVMS Operating System.

For more information regarding OpenVMS licensing terms and policies, contact your local VSI sales office, or find VSI software licensing information on the World Wide Web at: http://www.vmssoftware.com

License Management Facility Support

The OpenVMS operating system supports the OpenVMS License Management Facility (LMF). If an OpenVMS license is not registered and activated using LMF, only a single login is permitted for system management purposes through the system console (OPA0:).

OpenVMS compiler license types are based on the number of concurrent users, called an activity license. Every compiler product has the option to define an activity as related to the LMF. OpenVMS Interactive User and ADL Interactive User Licenses define the number of concurrent users who have activity licenses as defined by the LMF.

OpenVMS Alpha License Information

The following technologies are licensed as part of the OpenVMS Alpha operating system:

Product Name	Software Product Description (SPD)
DECnet-Plus for OpenVMS Alpha	SPD DO-VIBHAA-023
DECnet for OpenVMS Alpha	SPD DO-VIBHAA-024
DECram for OpenVMS	SPD DO-VIBHAB-005
DECwindows Motif for OpenVMS	SPD DO-VIBHAA-026
OpenVMS Cluster Software	SPD DO-VIBHAA-033
RMS Journaling for OpenVMS	SPD DO-VIBHAA-030
TCP/IP Services for OpenVMS	SPD DO-VIBHAA-029
Volume Shadowing for OpenVMS	SPD DO-VIBHAA-031

SYSTEMS SUPPORTED

AlphaServer Systems Supported

VSI OpenVMS Alpha Version 8.4-2L1 supports the following HPE AlphaServer systems:

- AlphaServer DS10
- AlphaServer DS15
- AlphaServer DS20
- AlphaServer DS20e
- AlphaServer DS25
- AlphaServer ES45
- AlphaServer ES47

Upgrading to VSI OpenVMS Alpha V8.4-2L1

VSI supports operating system upgrades only; fresh installations are not supported. VSI supports upgrades to VSI OpenVMS Alpha V8.4-2L1 from the following HPE OpenVMS versions:

- HPE Version v8.4 U900, U1000, U1100, U1200
- HPE Version v8.3
- HPE Version v7.3-2

APPENDIX A (VSI OpenVMS for AlphaServer systems)

This appendix describes the AlphaServer system options supported on VSI OpenVMS Alpha. Please refer to the Hewlett Packard Enterprise *QuickSpecs Online* for additional details regarding AlphaServer system disks, IO, SCSI, and FibreChannel options.

http://h41370.www4.hpe.com/quickspecs/overview.html

NOTE: The following items are not supported by VSI OpenVMS Alpha:

- ATM network protocol devices
- FDDI protocol devices
- Token ring LAN devices
- Memory channel adapters
- PCI to CI (CIPCA) adapters

VSI reserves the right to change the number and type of devices supported by OpenVMS Alpha, DECnet for OpenVMS Alpha, DECnet-Plus for OpenVMS, TCP/IP Services for OpenVMS, and OpenVMS Cluster software. The minimum hardware requirements for future versions and updates of these software products may be different from current hardware requirements. See the layered product SPDs for detailed product information about each product.

Terminals and Terminal Line Interfaces

To prevent input from overflowing a buffer, terminals use the ASCII control characters DC1, also known as XON, and DC3, also known as XOFF, for synchronization as defined by HP STD 111, Revision A. VXT windowing terminals support standard ANSI applications and X Windows Systems using the LAT transport protocol. OpenVMS Alpha supports the VT200 series, VT300 series, VT400 series, VT500 series, and VXT2000 series terminals.

Tapes

The following table lists tapes that are supported on OpenVMS Alpha V8.4-2L1, and the bus on which the device is supported.

Tape	Description	Bus
ESL9326	40/80 GB DLT tape library family	SCSI
ESL9198	40/80 GB DLT tape library family	SCSI
SDLT 320	320 GB SDLT Tape Drive	SCSI
SDLT 600	600 GB SDLT Tape Drive	SCSI
Ultrium 460	400 GB LTO Ultrium 2 Tape Drive	SCSI
Ultrium 448	400 GB LTO Ultrium 2 Tape Drive	SCSI
Ultrium 960	800 GB Ultrium 3 Tape Drive	SCSI/FC
Ultrium 920	800 LTO Ultrium 3 Tape Drive	SCSI/FC
Ultrium 1840	1.6 TB LTO Ultrium 4 Tape Drive	SCSI/FC
DAT72	72 GB DAT Tape Drive	SCSI
DAT160	160 GB DAT Tape Drive	SCSI
1/8 Autoloader	Tape Autoloader	SCSI
VLS6000	Virtual Tape Library	FC
MSL6000	Ultrium and SDLT Tape Library	FC
MSL2024	Ultrium Tape Library	SCSI/FC
MSL2048	Ultrium Tape Library	SCSI/FC
MSL8096	Ultrium Tape Library	SCSI/FC
ESL E-series	Ultrium and SDL Tape Library	FC
EML E-series	Ultrium Tape Library	FC

SAN-attached Tape Libraries

EML-E Series Tape Library	Enterprise Storage Library
ESL-E Series Tape Library	Enterprise Storage Library
ESL9595	Enterprise Storage Library
ESL9322	Enterprise Storage Library
ESL9326	Enterprise Storage Library
ESL9198	Enterprise Storage Library
MSL2024	Business Class Library
MSL4048	Business Class Library
MSL8096	Business Class Library
MSL5000 Series	Modular Storage Library
MSL6000 Series	Modular Storage Library
VLS 6000	Enterprise Virtual Tape Library
VLS9000	Virtual Tape Library
VLS12000	Virtual Tape Library
D2D4106	D2D backup systems
D2D4112	D2D backup systems
D2D4312	D2D backup systems
D2D4324	D2D backup systems

NOTE: OpenVMS supports both SDLT and Ultrium 460/960 tape drives within HP StorageWorks Tape libraries.

Networks Storage Servers

An Infoserver is an integrated hardware and software system that sits directly on the Ethernet to provide CD, hard disk, magneto-optical, and tape access to OpenVMS clients in a LAN. It supports up to 14 SCSI devices and can be used for software distribution and initial system load (ISL).

Enterprise Storage Arrays

EVA	StorageWorks Enterprise Virtual Array 8400, 8100, 8000, 6400, 6100, 6000, 5000, 4400, 4100, 4000, 3000, P6500, P6300
HSG60	Fibre Channel based StorageWorks controller that supports up to two UltraSCSI wide single-ended device ports and two host ports.
HSG80	Fibre Channel based StorageWorks controller that supports up to six UltraSCSI wide single-ended device ports and two host ports.
HSZ22	UltraSCSI wide differential based StorageWorks controller that supports up to two UltraSCSI wide single-ended device ports and two host ports.
HSZ80	UltraSCSI wide differential based StorageWorks controller that supports up to six UltraSCSI wide single-ended device ports and two host ports.
MSA	StorageWorks Modular Storage Array 1000, 1500
XP	StorageWorks XP Storage Array 128/1024, 48/512, P9500, 10000/12000, 20000/24000

Abbreviations

APMP	Adaptive Partitioned Multiprocessing
ATA	AT/Attachment
ATAPI	ATA Packet Interface
COM	Component Object Model
DLT	Digital Linear Tape
EISA	Extended Industry Standard Architecture
FSE	Fast Single Ended (SCSI)
FWD	Fast-Wide Differential (SCSI)
GigE	Gigabit Ethernet
IDE	Integrated Device (or Drive) Electronics
IEEE	Institute of Electrical and Electronics Engineers
I80277	Intel 82077 PC Compatible Floppy Interface
MSCP	Mass Storage Control Protocol
NCS	National Character Set
PCI	Peripheral Component Interconnect
QIC	Quarter Inch Cartridge
RAID	Redundant Array of Independent Disks
RMC	Remote Procedure Call
RMS	Record Management Services
SDI	Standard Drive Interface
SMP	Symmetric Multiprocessing
STI	Standard Tape Interface
TFF	Terminal Fallback Facility
TIE	Translated Image Environment
TMSCP	Tape Mass Storage Control Protocol
USB	Universal Serial Bus
VLM	Very Large Memory
XMI	Extended Memory Interconnect

SOFTWARE WARRANTY