



Software Product Description

PRODUCT NAME: OpenVMS Cluster Software

SPD 29.78.11

This Software Product Description describes Versions 6.2-1H1 and 7.0 of the following products:

- VMScluster Software for OpenVMS Alpha
- VAXcluster Software for OpenVMS VAX
- OpenVMS Cluster Client Software for Alpha (part of NAS150)
- OpenVMS Cluster Client Software for VAX (part of NAS150)

Except where noted features of Version 6.2-1H1 and 7.0 are identical.

Except where noted, the features described in this SPD apply equally to Alpha and VAX systems. OpenVMS Cluster Software licenses and part numbers are architecture specific; refer to the Ordering Information section of this SPD for further details.

DESCRIPTION

OpenVMS Cluster Software is an OpenVMS System Integrated Product (SIP). It provides a highly integrated OpenVMS computing environment distributed over multiple Alpha and VAX CPUs. In this SPD, this environment is referred to as a VMScluster.

CPUs in a VMScluster system can share processing, mass storage (including system disks), and other resources under a single OpenVMS security and management domain. Within this highly integrated environment, CPUs retain their independence because they use local, memory-resident copies of the OpenVMS operating system. Thus, VMScluster CPUs can boot and shut down independently while benefiting from common resources.

Applications running on one or more CPUs in a VMScluster system access shared resources in a coordinated manner. VMScluster software components synchronize access to shared resources, allowing multiple processes on any CPU in the VMScluster to perform coordinated, shared data updates.

Because resources are shared, VMScluster systems offer higher availability than standalone CPUs. Properly configured VMScluster systems can withstand the shut-down or failure of various components. For example, if one CPU in a VMScluster is shut down, users can log in to another CPU to create a new process and continue working. Because mass storage can be shared clusterwide, the new process is able to access the original data. Applications can be designed to survive these events automatically.

All VMScluster systems have the following software features in common:

- The OpenVMS operating system and VMScluster software allow all CPUs to share read and write access to disk files in a fully coordinated environment. Application programs can specify the level of clusterwide file sharing that is required; access is then coordinated by the OpenVMS Extended QIO Processor (XQP) and Record Management Services (RMS). Coherency of multi-CPU configurations is implemented by VMScluster software using a flexible and sophisticated per-CPU voting mechanism.
- Shared batch and print queues are accessible from any CPU in the VMScluster system. The OpenVMS queue manager controls clusterwide batch and print queues, which can be accessed by any CPU. Batch jobs submitted to clusterwide queues are routed to any available CPU so the batch load is shared.
- The OpenVMS Lock Manager System Services operate in a clusterwide manner. These services allow reliable coordinated access to any resource and provide signaling mechanisms at the system and process level across the whole VMScluster system.
- All physical disks and tapes in a VMScluster system can be made accessible to all CPUs.
- Process information and control services, including the ability to create and delete processes, are available on a clusterwide basis to application programs

and system utilities. (Clusterwide process creation is only available with Version 7.0.)

- Configuration command procedures assist in adding and removing CPUs and in modifying their configuration characteristics.
- The dynamic Show Cluster utility displays the status of VMScluster hardware components and communication links.
- A fully automated clusterwide data and application caching feature enhances system performance and reduces I/O activity.
- Standard OpenVMS system management and security features work in a clusterwide manner so that the entire VMScluster system operates as a single security and management domain.
- The VMScluster software dynamically balances the interconnect I/O load in VMScluster configurations that include multiple interconnects.
- Multiple VMScluster systems can be configured on a single or extended local area network (LAN). LANs and the LAN adapters used for VMScluster communications can be used concurrently by other network protocols.
- The optionally installable DECcmds availability management tool allows system managers to monitor and manage resource availability in real time on all the members of a VMScluster.
- Cross-architecture satellite booting permits VAX boot nodes to provide boot service to Alpha satellites and allows Alpha boot nodes to provide boot service to VAX satellites.
- System services are provided that enable applications to automatically detect changes in VMScluster membership.

Definitions

The following terms are used frequently throughout this SPD:

- Boot node — A CPU that is both a MOP server and a disk server. A boot node can fully service satellite boot requests.
- CPU (central processing unit) — An Alpha family or VAX family computer running the OpenVMS operating system. A CPU comprises one or more processors and operates as a VMScluster node. A VMScluster node can be referred to as VMScluster member.
- Disk server — A CPU that uses the OpenVMS MSCP server to make disks to which it has direct access available to other CPUs in the VMScluster system.
- HSC, HSJ — An intelligent mass storage controller subsystem that connects to the CI.
- HSD — An intelligent mass storage controller subsystem that connects to the DSSI.
- HSZ — An intelligent mass storage controller subsystem that connects to the SCSI.
- Maintenance Operations Protocol (MOP) server — A CPU that services satellite boot requests to provide the initial LAN downline load sequence of the OpenVMS operating system and VMScluster software. At the end of the initial downline load sequence, the satellite uses a disk server to perform the remainder of the OpenVMS booting process.
- Mixed-architecture VMScluster system — A VMScluster system that is configured with both VAX and Alpha CPUs.
- MSCP (mass storage control protocol) — A message-based protocol for controlling Digital Storage Architecture (DSA) disk storage subsystems. The protocol is implemented by the OpenVMS DUDRIVER device driver.
- Multihost — A configuration in which more than one CPU is connected to a single DSSI or SCSI bus.
- Satellite — A CPU that is booted over a LAN using a MOP server and disk server.
- Single-host — A configuration in which a single CPU is connected to a DSSI or SCSI bus.
- Star coupler — A common connection point for all CI connected CPUs and HSC and HSJ controllers.
- Tape server — A CPU that uses the OpenVMS TMSCP server to make tapes to which it has direct access available to other CPUs in the VMScluster system.
- TMSCP (tape mass storage control protocol) — A message-based protocol for controlling DSA tape-storage subsystems. The protocol is implemented by the OpenVMS TUDRIVER device driver.
- Vote — CPUs in a VMScluster system may be configured to provide Votes that are accumulated across the multi-CPU environment. Each CPU is provided with knowledge of how many votes are necessary to meet a quorum before distributed shared access to resources is enabled. A VMScluster system must be configured with at least one voting CPU.

OpenVMS Cluster Client Software

VMScluster configurations can be configured with CPUs that operate and are licensed explicitly as client systems. OpenVMS Cluster Client licensing is provided as part of the Digital NAS150 layered product package. OpenVMS Cluster Client CPUs contain full VMScluster functionality as described in this SPD, with the following exceptions:

- Client CPUs cannot provide votes toward the operation of the VMScluster system.
- Client CPUs cannot MSCP serve disks or TMSCP serve tapes.

Interconnects

VMScluster systems are configured by connecting multiple CPUs with a communications medium, referred to as an interconnect. VMScluster nodes communicate with each other using the most appropriate interconnect available. In the event of interconnect failure, VMScluster software automatically uses an alternate interconnect whenever possible. VMScluster software supports any combination of the following interconnects:

- CI (computer interconnect)
- DSSI (Digital Storage Systems Interconnect)
- SCSI (Small Computer Storage Interconnect)
- FDDI (Fiber Distributed Data Interface)
- Ethernet

CI and DSSI are highly optimized, special-purpose interconnects for CPUs and storage subsystems in VMScluster configurations.

SCSI is an industry-standard storage interconnect. Multiple CPUs can be configured on a single SCSI bus, thereby providing multihost access to SCSI storage devices. Note that the SCSI bus is not used for CPU-to-CPU communication. Consequently CPUs connected to a multihost SCSI bus must also be configured with another interconnect in order to provide CPU-to-CPU communication.

Ethernet and FDDI are industry-standard, general-purpose communications interconnects that can be used to implement a LAN. Except where noted, VMScluster support for both of these LAN types is identical.

VMScluster configurations can be configured using wide area networking (WAN) infrastructures, such as DS3 and ATM. Connection to these media is achieved with FDDI bridges or switches.

Configuration Rules

- The maximum number of CPUs supported in a VMScluster system is **96**.
- Every CPU in a VMScluster system must be connected to every other CPU via any of the supported VMScluster interconnects (see Table 1).
- VAX-11/7xx, VAX 6000, VAX 7000, VAX 8xxx, VAX 9000, and VAX 10000 series CPUs require a system disk that is accessed via a local controller or through a local CI or DSSI connection. These CPUs cannot be configured to boot as satellite nodes.
- All CPUs connected to a CI or DSSI must be configured as VMScluster members. VMScluster members configured on a CI or DSSI will become members of the same VMScluster (this is imposed automatically by the VMScluster software). All CPUs connected to a multihost SCSI bus must be configured as members of the same VMScluster.
- A VMScluster system can include any number of star couplers. The number of CI adapters supported by different CPUs can be found in Table 2. The number of star couplers that a CPU can be connected to is limited by the number of adapters with which it is configured.
- The maximum number of CPUs that can be connected to a star coupler is **16**, regardless of star coupler size.
- The KFQSA Q-bus to DSSI adapter does not support VMScluster communication to other CPUs on the DSSI; CPUs using this adapter must include another interconnect for VMScluster communication.
- The maximum number of CPUs that can be connected to a DSSI is **4**, regardless of system or adapter type. Any mix of systems and adapters is permitted. Depending on the CPU model, it may not be possible to configure four CPUs on a common DSSI bus because of DSSI bus cable-length restrictions. Refer to the specific CPU system configuration manuals for further information.
- The maximum number of CPUs that can be connected to a SCSI bus is **3**.
- The maximum number of multihost SCSI buses that a CPU may be connected to is **3**.
- VMScluster CPUs that are configured using WAN interconnects must adhere to the detailed line specifications described in the *OpenVMS New Features Manual*. The maximum CPU separation is 150 miles.
- A single time-zone setting must be used by all CPUs in a VMScluster system.

- A VMScluster system can be configured with a maximum of one quorum disk. A quorum disk cannot be a member of an OpenVMS volume set or of a shadow set created by the Volume Shadowing for OpenVMS product.
- A system disk can contain only a single version of the OpenVMS operating system and is architecture specific. For example, OpenVMS Alpha Version 7.0 *cannot coexist* on a system disk with OpenVMS VAX Version 7.0.
- HSJ and HSC series disks and tapes can be dual pathed between controllers on the *same or different* star couplers. The HSD30 series disks and tapes can be dual pathed between controllers on the *same or different* DSSI interconnects. Such dual pathing provides enhanced data availability using an OpenVMS automatic recovery capability called *failover*. Failover is the ability to use an alternate hardware path from a CPU to a storage device when a failure occurs on the current path. The failover process is transparent to applications. Dual pathing between an HSJ or HSC and a local controller is not permitted. When two local controllers are used for dual pathing, each controller must be located on a separate CPU of the same architecture. (Note: When disks and tapes are dual pathed between controllers connected to different star couplers or DSSI buses any CPU connected to one of the star couplers or buses must also be connected to the other.)
- Disks and tapes can be dual pathed between pairs of HSZ40 controllers that are arranged in a dual-redundant configuration. The controllers must be connected to the *same* host SCSI bus. Failover is accomplished using the HSZ40's transparent failover capability.
- OpenVMS operating system and layered-product installations and upgrades cannot be performed across architectures. OpenVMS Alpha software installations and upgrades must be performed using an Alpha system with direct access to its system disk. OpenVMS VAX software installations and upgrades must be performed using a VAX system with direct access to its system disk.
- Ethernet LANs and the protocols that use them must conform to the IEEE® 802.2 and IEEE 802.3 standards. Ethernet LANs must also support Ethernet Version 2.0 packet formats.
- FDDI LANs and the protocols that use them must conform to the IEEE 802.2, ANSI X3.139–1987, ANSI X3.148–1988, and ANSI X3.166–1990 standards.
- VMScluster systems support up to **4** LAN adapters per CPU for VMScluster communications.
- LAN segments can be bridged to form an extended LAN (ELAN). The ELAN must conform to IEEE® 802.1D, with the following restrictions:
 - All LAN paths used for VMScluster communication must operate with a nominal bandwidth of at least 10 megabits per second.
 - The ELAN must be capable of delivering packets that use the padded Ethernet Version 2.0 packet format and the FDDI SNAP/SAP packet format.
 - The ELAN must be able to deliver packets with a maximum data field length of at least 1080 bytes.¹
 - The maximum number of bridges between any two end nodes is **7**.
 - The maximum transit delay through any bridge must not exceed **2** seconds.
 - The ELAN must provide error-detection capability between end nodes that is equivalent to that provided by the Ethernet and FDDI data link frame-check sequences.
- The packet-retransmit timeout ratio for VMScluster traffic on the LAN from any CPU to another must be less than **1** timeout in **1000** transmissions.

Recommendations

The optimal VMScluster system configuration for any computing environment is based on requirements of cost, functionality, performance, capacity, and availability. Factors that impact these requirements include:

- Applications in use
- Number of users
- Number and models of CPUs
- Interconnect and adapter throughput and latency characteristics
- Disk and tape I/O capacity and access time
- Number of disks and tapes being served
- Interconnect utilization

Digital recommends VMScluster system configurations based on its experience with the VMScluster software product. The customer should evaluate specific application dependencies and performance requirements to determine an appropriate configuration for the desired computing environment.

¹ In the padded Ethernet format, the data field follows the two-byte length field. These two fields together comprise the LLC data field in the 802.3 format.

When planning a VMScluster system, consider the following recommendations:

- VMScluster CPUs should be configured using interconnects that provide appropriate performance for the required system usage. In general, use the highest performance interconnect possible. CI, DSSI, and FDDI are the preferred interconnects between powerful CPUs.
- Although VMScluster systems can include any number of system disks, consider system performance and management overhead in determining their number and location. While the performance of configurations with multiple system disks may be higher than with a single system disk, system management efforts increase in proportion to the number of system disks.
- Data availability and I/O performance are enhanced when multiple VMScluster nodes have direct access to shared storage; whenever possible, configure systems to allow direct access to shared storage in favor of OpenVMS MSCP served access. Multiaccess CI, DSSI, and SCSI storage provides higher data availability than singly accessed, local controller-based storage. Additionally, dual pathing of disks between local or HSC/H SJ/HSD/H SZ storage controllers enhances data availability in the event of controller failure.
- VMScluster systems can enhance availability by utilizing redundant components, such as additional CPUs, storage controllers, disks, and tapes. Extra peripheral options, such as printers and terminals, can also be included. Multiple instances of all VMScluster interconnects (CI, DSSI, SCSI, Ethernet, and FDDI) are supported.
- To enhance resource availability, VMSclusters that implement satellite booting should use multiple boot servers. When a server fails in configurations that include multiple servers, satellite access to multipath disks will fail over to another path. Disk servers should be the most powerful CPUs in the VMScluster and should use the highest bandwidth LAN adapters available.
- The performance of an FDDI LAN varies with each configuration. When an FDDI is used for VMScluster communications, the ring latency when the FDDI ring is idle should not exceed 400 microseconds. This ring latency translates to a cable distance between end nodes of approximately 40 kilometers.
- The ELAN must provide adequate bandwidth, reliability, and low delay in order to optimize the operation of the VMScluster. The average LAN segment utilization should not exceed 60% for any 10-second interval. If ELAN performance degrades to the point where nodes cannot communicate every 3 seconds, then nodes may leave the VMScluster. The effective performance of the ELAN can be increased by following these guidelines:
 - Configure high-performance nodes with multiple LAN adapters connected to different LAN segments.
 - Minimize the number of bridges on the path between nodes that communicate frequently, such as satellites and their boot servers.
 - Use bridges to isolate and localize the traffic between nodes that communicate with each other frequently. For example, use bridges to separate the VMScluster from the rest of the ELAN and to separate nodes within a cluster that communicate frequently from the rest of the VMScluster.
 - Use FDDI on the communication paths that have the highest performance requirements. The NISCS_MAX_PKTSZ system parameter can be adjusted to use the full FDDI packet size. Ensure that the ELAN path supports a data field of at least 4470 bytes end to end, or the ELAN path sets the priority field to zero in the FDDI frame-control byte on the destination FDDI link.
 - Minimize the packet delay between end nodes.
- The RAID level 1 storage functionality of Volume Shadowing for OpenVMS provides the following advantages:
 - Enhanced data availability in the event of disk failure
 - Enhanced read performance with multiple shadow-set members

For more information, refer to the Volume Shadowing for OpenVMS Software Product Description (SPD 27.29.xx).
- The DECram for OpenVMS software product can be used to create very high-performance, memory-resident RAM disks. Refer to the DECram for OpenVMS Software Product Description (SPD 34.26.xx) for additional information.

DECamsd Features

VMScluster software incorporates the features of a real-time monitoring, investigation, diagnostic, and system management tool that can be used to improve system availability.

The DECamsd availability management tool contains a console and an OpenVMS device driver. The console is a DECwindows Motif® based application that allows system managers to display windows showing processes, quotas, disks, locks, memory, SCS data structures, and I/O activity in the VMScluster. The Motif display can be directed to any X-compatible display. The driver is a data collector that runs on the monitored VM-Cluster members. Console application and driver software is provided for Alpha and VAX systems.

HARDWARE SUPPORT*CPU Support*

Any Alpha or VAX CPU, as documented in the OpenVMS Operating System for VAX and Alpha Software Product Description (SPD 25.01.xx), can be used in a VMScluster.

Peripheral Option and Storage Controller Support

VMScluster systems can use all peripheral options and storage subsystems supported by OpenVMS. Refer to the OpenVMS Operating System for VAX and Alpha SPD for more information.

Interconnect Support

Table 1 shows which processors are supported on which interconnects and whether the processor can be booted as a satellite node over that interconnect. All CPUs can service satellite boot requests over a LAN interconnect (FDDI or Ethernet).

Note: Levels of interconnect support and LAN booting capabilities are continuously being increased. In many cases these additional capabilities result from hardware option and system console microcode enhancements and are not dependent on OpenVMS software. Refer to the appropriate hardware option and system documentation for the most up-to-date information.

LAN Support

VMScluster systems can use all Ethernet and FDDI LAN adapters supported by OpenVMS for access to Ethernet and FDDI interconnects. Refer to the OpenVMS Operating System for VAX and Alpha SPD for more information.

Table 1

CPU	CI	DSSI	SCSI ¹	FDDI	Ethernet
AlphaServer 8200 ² , 8400	Yes ³	Yes ⁴	–	Yes+Sat ^{5, 6}	Yes
DEC 7000, 10000	Yes	Yes	–	Yes+Sat ⁶	Yes
DEC 4000	–	Yes	–	Yes	Yes+Sat
DEC 3000	–	–	–	Yes+Sat ⁷	Yes+Sat
AlphaServer 2100, 2100A	–	Yes	Yes	Yes+Sat ⁸	Yes+Sat
AlphaServer 1000, 2000	–	Yes	Yes	Yes	Yes+Sat
AlphaServer 400	–	–	Yes	Yes	Yes+Sat
AlphaStation 200, 250, 400, 600	–	–	Yes	Yes	Yes+Sat
DEC 2000	–	–	–	Yes	Yes+Sat
VAX 6000, 7000, 10000	Yes	Yes	–	Yes	Yes
VAX 8xxx, 9xxx, 11/xxx	Yes	–	–	–	Yes
VAX 4xxx ⁹	–	Yes	–	Yes	Yes+Sat
VAX 2xxx, 3xxx ⁹	–	–	–	–	Yes+Sat

¹This column refers to multihost SCSI connectivity. Refer to the appropriate system documentation for information regarding single-host connectivity to SCSI buses.

²Rack-mount AlphaServer 8200 configurations with CIXCD, KFMSB, and DEMFA capability are available from Digital CSS.

³Each "Yes" means that this CPU is supported on this interconnect but cannot be booted as a satellite over this interconnect.

⁴Connection to the DSSI is achieved using the KFMSB. These systems do not support the KFESA/B.

⁵Each "Yes+Sat" means that this CPU is supported on this interconnect and can be booted as a satellite node over this interconnect.

⁶Using DEMFA only.

⁷Using DEFTA only.

⁸V7.0: using DEFEA only. V6.2-1H1: using DEFEA/DEPFA only.

⁹Some models may provide slightly different interconnect support. Refer to the system specific hardware manual for complete details.

CI Support

VMScluster nodes can be configured with multiple CI adapters. Table 2 shows the types of adapters that are supported by each CPU. There can be only one type of adapter configured on a CPU; the maximum quantity of each type is noted in the table. The CI adapters in a CPU can connect to the same or different star couplers.

Note: The CIBCA-A adapter cannot coexist with a KFMSA adapter on the same CPU.

Note: The CIBCA-A and CIBCA-B are different.

Table 2

CPU Type	CI750	CI780	CIBCI	CIBCA-A	CIBCA-B	CIXCD
AlphaServer 8400	-	-	-	-	-	10
DEC 7000, 10000	-	-	-	-	-	10
VAX 11/750	1	-	-	-	-	-
VAX 11/780, 11785	-	1	-	-	-	-
VAX 6000	-	-	-	1	4	4
VAX 82xx, 83xx	-	-	1	1	1	-
VAX 86xx	-	2	-	-	-	-
VAX 85xx, 8700, 88xx	-	-	1	1	2	-
VAX 9000	-	-	-	-	-	6
VAX 7000, 10000	-	-	-	-	-	10

Star Coupler Expander

A CI star coupler expander (CISCE) can be added to any star coupler to increase its connection capacity to 32 ports. The maximum number of CPUs that can be connected to a star coupler is **16**, regardless of the number of ports.

The DEFZA FDDI adapter is supported on VAX systems only.

DSSI Support

Any mix of Alpha and VAX DSSI adapters can be configured on a common DSSI bus. Refer to the appropriate hardware manuals for specific adapter and configuration information. The following points provide general guidelines for configurations:

- Configure VAX 6000, VAX 7000, and VAX 10000 systems with KFMSA adapters.
- Configure DEC 7000, DEC 10000, and AlphaServer 8400/8200 XMI systems with KFMSB adapters.
- Up to **6** KFMSA/Bs can be configured on an XMI bus. Up to **12** KFMSA/Bs can be configured in a system.
- Configure the AlphaServer systems shown in Table 1 with KFESB adapters. However, note that AlphaServer 8200/8400 do not support the KFESB adapter, and should be configured with the KFMSB adapter. The AlphaServer 2100 can also be configured with KFESA adapters. AlphaStation systems cannot be configured with KFESA/B.
- Up to **4** KFESBs can be configured on a system. Up to **2** KFESAs can be configured on a system. A mix of **1** KFESB and **1** KFESA can be configured on a system.

- Because the DEC 4000 DSSI adapter terminates the DSSI bus, only two DEC 4000s can be configured on a DSSI.

Multihost SCSI Support

OpenVMS Cluster Software provides support for multihost SCSI configurations using a restricted range of Alpha systems and SCSI adapters, devices, and controllers. Single-host SCSI support is provided for an extensive range of systems and SCSI adapters, devices and controllers. For further information about the complete range of SCSI support, refer to the OpenVMS Operating System for VAX and Alpha SPD.

Table 1 shows which systems can be configured on a multihost SCSI bus.

For Version 7.0, these systems must use their embedded system SCSI adapters or optional KZPAA (single-ended) adapters to connect to a multihost SCSI bus.

For Version 6.2-1H1, AlphaServer 2000, 2100 and 2100A systems may use either their embedded system SCSI adapters or optional KZPAA or KZPSA (fast wide differential) adapters to connect to a multihost SCSI bus.

Note: Any supported SCSI adapter can be used to connect to a single-host SCSI bus.

Optional adapters are recommended for connection to multihost buses. Use of optional adapters simplifies SCSI cabling and also leaves the embedded system adapter available for tape drives, floppies, and CD-ROMs.

Multihost SCSI configurations can include DWZZA /DWZZB single-ended SCSI to differential SCSI converters.

Multihost SCSI buses can be configured with any appropriately compliant SCSI-2 disk. Disks must support the following three features:

- Multihost support
- Tagged command queuing
- Automatic bad block revectoring

SCSI disk requirements are fully documented in the the *OpenVMS Version 6.2-1H1 and 7.0 New Features* manuals. As a guide, the following storage devices have been qualified by Digital for use on multihost SCSI buses. Both wide and narrow variants of these disks have been qualified. (Note that the KZPAA is a narrow adapter and does not enable wide-mode operation.)

- RZ28
- RZ28B
- RZ28M
- RZ26

- RZ26L
- RZ26N
- RZ29B

Note: RZ25 disks do not support tagged command queuing and therefore are not supported for use on multihost SCSI buses.

Tape drives, floppies, and CD-ROMs cannot be configured on multihost SCSI buses. Configure these devices on single-host SCSI buses.

Multihost SCSI buses must adhere to all SCSI-2 specifications. Rules regarding cable length and termination must be adhered to carefully. Refer to the SCSI-2 specification or the *OpenVMS Version 6.2-1H1 and 7.0 New Features* manuals for further information.

DECamsds Console

Digital recommends that the availability management console run on a standalone workstation with a color monitor. However, it can also run on a workstation that is configured as a VMScluster member or on a nonworkstation system using DECwindows to direct the display to an X-based display.

SOFTWARE REQUIREMENTS

- OpenVMS Operating System

Refer to the OpenVMS Operating System for VAX and Alpha Software Product Description (SPD 25.01.xx) for more information.

The ability to have more than one version of OpenVMS in a VMScluster allows upgrades to be performed in a staged fashion so that continuous VMScluster system operation is maintained during the upgrade process. Only one version of OpenVMS can exist on any system disk; multiple versions of OpenVMS in a VMScluster require multiple system disks. Also note that system disks are architecture specific: OpenVMS Alpha and OpenVMS VAX cannot coexist on the same system disk. The coexistence of multiple versions of OpenVMS in a VMScluster configuration is supported according to the following conditions:

- *Warranted support* is provided for mixed-architecture VMSclusters in which all Alpha and VAX systems are running OpenVMS Version 7.0.
- *Warranted support* is provided for mixed-architecture VMSclusters in which all Alpha and VAX systems are running any variant of OpenVMS Version 6.2.

Warranted support means that Digital has fully qualified the two architectures coexisting in a VMScluster and will answer any problems identified by customers using these configurations.

— *Migration support* is provided for VMSclusters running any variant of OpenVMS Version 6.2 and OpenVMS Versions 1.5 (Alpha), 5.5-2 and 6.0 (VAX), and 6.1 (Alpha and VAX).

— *Migration support* is provided for VMSclusters running OpenVMS Version 7.0 and any variant of OpenVMS Version 6.2. Coexistence of OpenVMS Version 7.0 with releases prior to OpenVMS Version 6.2 is not supported but in most cases will operate satisfactorily.

Migration support means that Digital has qualified the two architectures and versions for use together in configurations that are migrating in a staged fashion to a higher version of OpenVMS or to Alpha systems. Digital will answer problem reports submitted about these configurations. However, in exceptional cases, Digital may recommend that you move your system to a warranted configuration as part of the solution.

Note that Digital does not support the use of more than two versions of OpenVMS software in a VMScluster at a time.

Digital recommends that all Alpha and VAX systems in a VMScluster run the latest version of OpenVMS.

- DECnet software

DECnet software is not required in a VMScluster configuration. However, DECnet software is necessary if internode process to process communication using DECnet mailboxes is needed.

The OpenVMS Version 6.2-1H1 Monitor utility uses DECnet for intra-cluster communication.

The OpenVMS Version 7.0 Monitor utility uses TCP/IP or DECnet based transports, as appropriate, for intra-cluster communication.

Refer to the appropriate DECnet Software Product Description for further information.

- DECamsds Availability Manager

The DECamsds Availability Manager requires DECwindows Motif for OpenVMS (SPD 42.19.xx).

OPTIONAL SOFTWARE

For information about VMScluster support for optional software products, refer to the *VMScluster Support* section of the Software Product Descriptions for those products.

Optional products that may be useful in VMScluster systems include:

- Volume Shadowing for OpenVMS (SPD 27.29.xx)
- StorageWorks RAID Software for OpenVMS (SPD 46.49.xx)

- DECram for OpenVMS (SPD 34.26.xx)
- POLYCENTER Performance Data Collector for OpenVMS (SPD 36.02.xx)
- POLYCENTER Performance Advisor for OpenVMS (SPD 36.03.xx)
- VAXcluster Console System (SPD 27.46.xx)
- Business Recovery Server (SPD 35.05.xx)

GROWTH CONSIDERATIONS

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

DISTRIBUTION MEDIA

OpenVMS Cluster Software is distributed on the same distribution media as the OpenVMS Operating System. Refer to the OpenVMS Operating System for VAX and Alpha SPD for more information.

ORDERING INFORMATION

OpenVMS Cluster Software is orderable as follows:

Every server (nonclient) Alpha system in a VMScluster configuration requires:

- VMScluster Software for OpenVMS Alpha
 - Software Licenses: QL-MUZA*-AA
 - Software Product Services: QT-MUZA*-**

Every server (nonclient) VAX system in a VMScluster configuration requires:

- VAXcluster Software for OpenVMS VAX
 - Software Licenses: QL-VBRA*-AA
 - Software Product Services: QT-VBRA*-**

OpenVMS Cluster Client Software is available as part of the NAS150 product. It is not separately orderable.

* Denotes variant fields. For additional information on available licenses, services, and media, refer to the appropriate price book.

The right to the functionality of the OpenVMS Cluster Availability Manager (DECamds) is included in all the licenses in the preceding list.

DOCUMENTATION

The *VMScluster Systems for OpenVMS* manual, the *Guidelines for VMScluster Configurations* manual, and the *DECamds User's Guide* are included in the OpenVMS hardcopy documentation as part of the Full Documentation Set.

Refer to the OpenVMS Operating System for VAX and Alpha Software Product Description for additional information about OpenVMS documentation and ordering information.

Specific terms and conditions regarding documentation on media apply to this product. Refer to Digital's terms and conditions of sale, as follows:

"A software license provides the right to read and print software documentation files provided with the software distribution kit for use by the licensee as reasonably required for licensed use of the software. Any hard copies or copies of files generated by the licensee must include Digital's copyright notice. Customization or modifications, of any kind, to the software documentation files are not permitted.

Copies of the software documentation files, either hardcopy or machine readable, may only be transferred to another party in conjunction with an approved relicense by Digital of the software to which they relate."

SOFTWARE LICENSING

This software is furnished under the licensing provisions of Digital Equipment Corporation's Standard Terms and Conditions. For more information about Digital's licensing terms and policies, contact your local Digital office.

License Management Facility Support

The OpenVMS Cluster Software product supports the OpenVMS License Management Facility (LMF).

License units for this product are allocated on an Unlimited System Use basis.

For more information about the License Management Facility, refer to the OpenVMS Operating System for VAX and Alpha Software Product Description (SPD 25.01.xx) or documentation set.

SOFTWARE PRODUCT SERVICES

A variety of service options are available from Digital. For more information, contact your local Digital office.

SOFTWARE WARRANTY

Warranty for this software product is provided by Digital with the purchase of a license for the product as defined in the Software Warranty Addendum of this SPD.

The above information is valid at time of release. Contact your local Digital office for the most up-to-date information.

© 1995 Digital Equipment Corporation.
All rights reserved.

™ AlphaServer, AlphaStation, BI, Business Recovery Server, CI, DECams, DECchip, DECnet, DECram, DECwindows, DELUA, DEUNA, Digital, DSSI, HSC, HSC40, HSC50, HSC60, HSC70, HSC90, HSJ, HSZ, MicroVAX, MicroVAX II, MSCP, OpenVMS, POLYCENTER, Q-bus, RA, RZ, StorageWorks, TA, TMSCP, TURBOchannel, UNIBUS, VAX, VAX 6000, VAX 9000, VAX-11/750, VAX-11/780, VAXstation, VAXcluster, VMScluster XMI, and the DIGITAL logo are trademarks of Digital Equipment Corporation.

IEEE is a registered trademark of the Institute of Electrical and Electronics Engineers, Inc.

Motif is a registered trademark of the Open Software Foundation, Inc.