Software Product Description

PRODUCT NAME: HP OpenVMS Cluster Software **SPD 29.78.24**

This Software Product Description describes Versions 6.2–1H3, 7.1–1H1, 7.1–1H2, 7.1–2, 7.2, 7.2–1, 7.2–1H1, 7.2-2, 7.3, 7.3–1, and 7.3–2 of the following products:

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

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 or VAX systems, or a mix of Alpha and VAX systems. In this SPD, this environment is referred to as an OpenVMS Cluster.

Systems in an OpenVMS Cluster 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, systems retain their independence because they use local, memory-resident copies of the OpenVMS operating system. Thus, OpenVMS Cluster systems can boot and shut down independently while benefiting from common resources.

Applications running on one or more systems in an OpenVMS Cluster system can access shared resources in a coordinated manner. OpenVMS Cluster software components synchronize access to shared resources, allowing multiple processes on any system in the OpenVMS Cluster to perform coordinated, shared data updates.

Because resources are shared, OpenVMS Cluster systems offer higher availability than standalone systems. Properly configured OpenVMS Cluster systems can withstand the shutdown or failure of various components. For example, if one system in an OpenVMS Cluster is shut down, users can log in to another system 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 OpenVMS Cluster systems have the following software features in common:

- The OpenVMS operating system and OpenVMS Cluster software allow all systems 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 multiple-system configurations is implemented by OpenVMS Cluster software using a flexible and sophisticated per-system voting mechanism.
- Shared batch and print queues are accessible from any system in the OpenVMS Cluster system. The OpenVMS queue manager controls clusterwide batch and print queues, which can be accessed by

any system. Batch jobs submitted to clusterwide queues are routed to any available system 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 OpenVMS Cluster system.
- All disks and tapes in an OpenVMS Cluster system can be made accessible to all systems.
- 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 available with Version 7.1 and higher.)
- Configuration command procedures assist in adding and removing systems and in modifying their configuration characteristics.
- The dynamic Show Cluster utility displays the status of OpenVMS Cluster hardware components and communication links.
- A fully automated clusterwide data and application caching feature enhances system performance and reduces I/O activity.
- The ability to define logical names that are visible across multiple nodes in an OpenVMS Cluster (Version 7.2 and higher).
- An application programming interface (API) allows applications within multiple OpenVMS Cluster nodes to communicate with each other (Version 7.2 and higher).
- Standard OpenVMS system management and security features work in a clusterwide manner so that the entire OpenVMS Cluster system operates as a single security and management domain.
- The OpenVMS Cluster software dynamically balances the interconnect I/O load in OpenVMS Cluster configurations that include multiple interconnects.
- Multiple OpenVMS Cluster systems can be configured on a single or extended local area network (LAN). LANs and the LAN adapters used for OpenVMS Cluster communications can be used concurrently by other network protocols.
- The optionally installable DECamds availability management tool (as well as HP Availability Manager) allows system managers to monitor and manage resource availability in real time on all the members of an OpenVMS Cluster.
- 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 enable applications to automatically detect changes in OpenVMS Cluster membership.

Definitions

The following terms are used frequently throughout this SPD:

- Boot node — A system that is both a MOP server and a disk server. A boot node can fully service satellite boot requests.
- System — An Alpha family or VAX family computer running the OpenVMS operating system. A system comprises one or more processors and operates as an OpenVMS Cluster node. An OpenVMS Cluster node can be referred to as an OpenVMS Cluster member.
- Disk server — A system that uses the OpenVMS MSCP server to make disks to which it has direct access available to other systems in the OpenVMS Cluster system.
- HSC, HSJ — An intelligent mass storage controller subsystem that connects to the CI bus.
- HSD — An intelligent mass storage controller subsystem that connects to the DSSI bus.
- HSG, HSV, MSA, XP — An intelligent mass storage controller subsystem that connects to the Fibre Channel bus.
- HSZ — An intelligent mass storage controller subsystem that connects to the SCSI bus.
- MDR (Modular Data Router) — Fibre Channel to SCSI bridge allowing SCSI tape devices to be used behind a Fibre Channel switch.
- NSR (Network Storage Router) — Fibre Channel to SCSI bridge allowing SCSI tape devices to be used behind a Fibre Channel switch.
- Maintenance Operations Protocol (MOP) server — A system that services satellite boot requests to provide the initial LAN downline load sequence of the OpenVMS operating system and OpenVMS Cluster 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 OpenVMS Cluster system — An OpenVMS Cluster system that is configured with both VAX and Alpha systems.

- 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 configuration — A configuration in which more than one system is connected to a single CI, DSSI, SCSI, or Fibre Channel interconnect.
- Satellite — A system that is booted over a LAN using a MOP server and disk server.
- Single-host configuration — A configuration in which a single system is connected to a CI, DSSI, SCSI, or Fibre Channel interconnect.
- Star coupler — A common connection point for all CI connected systems and HSC and HSJ controllers.
- Tape server — A system that uses the OpenVMS TMSCP server to make tapes to which it has direct access available to other systems in the OpenVMS Cluster 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 — Systems in an OpenVMS Cluster system can be configured to provide votes that are accumulated across the multi-system environment. Each system is provided with knowledge of how many votes are necessary to meet a quorum before distributed shared access to resources is enabled. An OpenVMS Cluster system must be configured with at least one voting system.

OpenVMS Cluster Client Software

OpenVMS Cluster configurations can be configured with systems that operate and are licensed explicitly as client systems. OpenVMS Cluster Client licensing is provided as part of the NAS150 layered product. An individually available license for DS-series AlphaServers is also provided. OpenVMS Cluster Client systems contain full OpenVMS Cluster functionality as described in this SPD, with the following exceptions:

- Client systems cannot provide votes toward the operation of the OpenVMS Cluster system.
- Client systems cannot MSCP serve disks or TMSCP serve tapes.

Interconnects

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

- CI (computer interconnect)
- DSSI (Digital Storage Systems Interconnect)
- SCSI (Small Computer Storage Interconnect) (storage only, Alpha only)
- FDDI (Fiber Distributed Data Interface)
- Ethernet (10/100, Gigabit)
- Asynchronous transfer mode (ATM) (emulated LAN configurations only)
- Memory Channel (Version 7.1 and higher only)
- Fibre Channel (storage only, Alpha only, Version 7.2–1 and higher only)

CI and DSSI are highly optimized, special-purpose interconnects for systems and storage subsystems in OpenVMS Cluster configurations. CI and DSSI provide both system-to-storage communication and system-to-system communication.

SCSI is an industry-standard storage interconnect. Multiple systems 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 system-to-system communication. Consequently, systems connected to a multihost SCSI bus must also be configured with another interconnect to provide system-to-system communication.

Fibre Channel is an evolving industry standard interconnect for storage and communications. Support by OpenVMS Version 7.2-1 (and higher) allows for a storage-only interconnect in a multihost environment utilizing Fibre Channel switched topologies. With Version 7.2-2, support for SCSI Tapes utilizing the Modular Data Router bridge, and the Network Storage Router bridge, is supported. As is true with SCSI, systems connected to a multihost Fibre Channel bus must also be configured with another interconnect to provide system-to-system communication.

Ethernet, ATM, and FDDI are industry-standard, general-purpose communications interconnects that can be used to implement a local area network (LAN). Except where noted, OpenVMS Cluster support for these LAN types is identical. The ATM device must be used as an emulated LAN configured device. Ethernet and FDDI provide system-to-system communication. Storage can be configured in FDDI environments that support FDDI-based storage servers.

OpenVMS Cluster configurations can be configured using wide area network (WAN) infrastructures, such as DS3, E3, and ATM. Connection to these media is achieved by the use of WAN interswitch links (ISLs).

Memory Channel is a high-performance interconnect that provides system-to-system communication. Memory Channel does not provide direct access to storage, so a separate storage interconnect is required in Memory Channel configurations.

Configuration Rules

- The maximum number of systems supported in an OpenVMS Cluster system is **96**.
- Every system in an OpenVMS Cluster system must be connected to every other system via any supported OpenVMS Cluster interconnect (see Table 1).
- VAX-11/7xx, VAX 6000, VAX 7000, VAX 8xxx, VAX 9000, and VAX 10000 series systems require a system disk that is accessed via a local adapter or through a local CI or DSSI connection. These systems cannot be configured to boot as satellite nodes.
- All systems connected to a common CI, DSSI, or Memory Channel interconnect must be configured as OpenVMS Cluster members. OpenVMS Cluster members configured on a CI, DSSI, or Memory Channel will become members of the same OpenVMS Cluster (this is imposed automatically by the OpenVMS Cluster software). All systems connected to a multihost SCSI bus must be configured as members of the same OpenVMS Cluster.
- An OpenVMS Cluster system can include any number of star couplers. Table 2 shows the number of CI adapters supported by different systems. The number of star couplers that a system can be connected to is limited by the number of adapters with which it is configured.
- The maximum number of systems 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 system-to-system communication across the DSSI; systems using this adapter must include another interconnect for system-to-system communication.
- The maximum number of systems that can be connected to a DSSI is four, regardless of system or adapter type. Any mix of systems and adapters is permitted, except where noted in the Hardware Support section of this SPD. Depending on the system model, it may not be possible to configure four systems on a common DSSI bus because of DSSI bus cable-length restrictions. Refer to the specific system system configuration manuals for further information.
- The maximum number of systems that can be connected to a SCSI bus is **3**. If the SCSI bus includes a five-port or greater Fair Arbitration SCSI Hub (DWZZH–05), the maximum number of systems is increased to **4**.
- The maximum number of multihost SCSI buses that a system can be connected to is **26**.
- The configuration size for Fibre Channel storage increases on a regular basis with new updates to OpenVMS. As such, please refer to the *Guidelines for OpenVMS Cluster Configurations* manual for the most up-to-date configuration capabilities.
- Beginning with OpenVMS Version 7.2–1, Multipath Failover for both Parallel SCSI and Fibre Channel storage environments is supported. This feature allows for the failover of cluster storage communications from one path to another when multiple storage buses have been connected to the same data source. For detailed information, refer to the *Guidelines for OpenVMS Cluster Configurations* manual.
- Beginning with OpenVMS Version 7.3–1, Multipath failover to the MSCP served path is supported. This feature allows failovers from physical connected storage paths to the cluster served path for data access. For detailed information, refer to the *Guidelines for OpenVMS Cluster Configurations* manual.
- OpenVMS Cluster systems that are configured using WAN interconnects must adhere to the detailed line specifications described in the *Guidelines for OpenVMS Cluster Configurations* manual. The maximum system separation is 150 miles.
- A single time-zone setting must be used by all systems in an OpenVMS Cluster system.
- An OpenVMS Cluster 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 HP 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 7.3–2 *cannot* coexist on a system disk with OpenVMS VAX Version 7.3.
- 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 system 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 adapter is not permitted. When two local adapters are used for dual pathing, each adapter must be located on a separate system of the same architecture. (Note: When disks and tapes are dual pathed between controllers that are connected to different star couplers or DSSI buses, any system connected to one of the star couplers or buses must also be connected to the other.)

- Disks can be dual pathed between pairs of HSZ 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 HSZ 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.
- 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 OpenVMS Cluster 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 average packet-retransmit timeout ratio for OpenVMS Cluster traffic on the LAN from any system to another must be less than **1** timeout in **1000** transmissions.

Recommendations

The optimal OpenVMS Cluster 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 systems
- 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

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

When planning an OpenVMS Cluster system, consider the following recommendations:

- OpenVMS Cluster systems should be configured using interconnects that provide appropriate performance for the required system usage. In general, use the highest-performance interconnect possible. Gigabit Ethernet and Memory Channel are the preferred interconnects between powerful systems.
- Although OpenVMS Cluster 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 OpenVMS Cluster systems 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, SCSI, and Fibre Channel storage provides higher data availability than singly accessed, local adapter-based storage. Additionally, dual pathing of disks between local

[1] In the padded Ethernet format, the data field follows the 2-byte length field. These two fields together comprise the LLC data field in the 802.3 format.

or HSC/HSJ/HSD/HSZ/HSG storage controllers enhances data availability in the event of controller failure.

- OpenVMS Cluster systems can enhance availability by utilizing redundant components, such as additional systems, storage controllers, disks, and tapes. Extra peripheral options, such as printers and terminals, can also be included. Multiple instances of all OpenVMS Cluster interconnects (CI, Memory Channel, DSSI, Ethernet, ATM, Gigabit Ethernet, FDDI, and of all OpenVMS Cluster Storage interconnects (SCSI and Fibre Channel) are supported.
- To enhance resource availability, OpenVMS Clusters 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 systems in the OpenVMS Cluster 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 OpenVMS Cluster 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 to optimize the operation of the OpenVMS Cluster. There are in-depth configuration guidelines for these ELAN environments in the OpenVMS documentation set, which are frequently updated as the technology area evolves. For specific configuration information, refer to the following manuals:
  - — *OpenVMS Cluster Systems*
  - — *Guidelines for OpenVMS Cluster Configurations*
- 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 shadowset members

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

OpenVMS Cluster Management Tools

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

HP DECamds

The HP DECamds 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 OpenVMS Cluster. The Motif™ display can be directed to any X-compatible display. The driver is a data collector that runs on the monitored OpenVMS systems. Console application and driver software is provided for Alpha and VAX systems.

HP Availability Manager

HP Availability Manager is functionally similar to DECamds, but it runs on Windows-based systems and on OpenVMS Alpha.

SCACP

Systems Communications Architecture Control Program (SCACP) is designed to monitor and manage LAN cluster communications.

HARDWARE SUPPORT

System Support

Any Alpha or VAX system, as documented in the *HP OpenVMS Operating System for VAX and Alpha Software Product Description* (SPD 25.01.xx), can be used in an OpenVMS Cluster.

Peripheral Option and Storage Controller Support

OpenVMS Cluster systems can use all peripheral options and storage subsystems supported by OpenVMS. Refer to the *OpenVMS Operating System for VAX and Alpha SPD* (SPD 25.01.xx) for more information.

Interconnect Support

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

Note: Levels of interconnect support and LAN booting capabilities are continually being increased. In many cases, these additional capabilities result from hardware option and system console microcode enhancements and are not dependent on OpenVMS software. For

the most up-to-date information, refer to the appropriate hardware option and system documentation.

LAN Support

OpenVMS Cluster systems can use all Ethernet (10 Mb/sec and 100 Mb/sec) and FDDI LAN adapters supported by OpenVMS for access to Ethernet and FDDI interconnects. Any number of LAN adapters can be configured in any combination (with the exception that a Q-bus can be configured with only one FDDI adapter). Refer to the *OpenVMS Operating System for VAX and Alpha Software Product Description* (SPD 25.01.xx) for more information.

Gigabit Ethernet LAN adapters can be used for *limited* OpenVMS Cluster interconnect capability for Version 7.1-2 through Version 7.2-*xx*. OpenVMS Version 7.3 and higher clusters provide more robust support for Gigabit Ethernet and ATM emulated LAN Ethernet connections. Additionally, OpenVMS Version 7.3 and higher also allows for load distribution of SCS cluster communications traffic across multiple, parallel LAN connections between cluster nodes. For specific limitations on these interconnects, refer to the release notes for your OpenVMS operating system version.

The DEFZA FDDI adapter is supported on VAX systems only.

Note: VAX systems cannot be booted over an FDDI.

Table 1

System	CI	Memory Channel[1]	DSSI	Multi-Host SCSI	FDDI	ATM,[3] Ethernet	Fibre Channel
AlphaServer GS 80/160/320, GS60/140, GS1280, 8200, 8400	Yes[4]	Yes	Yes[5]	Yes	Yes+Sat[6]	Yes	Yes
AlphaServer ES40, ES45, ES47, ES80, 4000, 4100	Yes	Yes	Yes	Yes	Yes+Sat	Yes	Yes
AlphaServer 1200, 2000, 2100, 2100A, DS20, DS20E	Yes	Yes	Yes	Yes	Yes+Sat	Yes+Sat	Yes[7]
AlphaServer DS10/10L/20/25, 1000, 1000A	–	Yes	Yes	Yes	Yes+Sat	Yes+Sat	Yes[8]
AlphaServer 400,800	–	–	Yes	Yes	Yes+Sat[1]	Yes+Sat	Yes[9]
AlphaServer 300	–	–	–	Yes	Yes	Yes+Sat	–
AlphaStations	–	–	–	Yes	Yes+Sat[10]	Yes+Sat	Yes[2]
DEC 7000, 10000	Yes	–	Yes	–	Yes+Sat	Yes	–
DEC 4000	–	–	Yes	–	Yes	Yes+Sat	–
DEC 3000	–	–	–	Yes	Yes+Sat[11]	Yes+Sat	–
DEC 2000	–	–	–	–	Yes	Yes+Sat	–
VAX 6000, 7000, 10000	Yes	–	Yes	–	Yes	Yes	–
VAX 8xxx, 9xxx, 11/xxx	Yes	–	–	–	–	Yes	–
VAX 4xxx[12]	–	–	Yes	–	Yes	Yes+Sat	–
VAX 2xxx, 3xxx[12]	–	–	–	–	–	Yes+Sat	–

[1] Version 7.1 and higher only. Support for Memory Channel on the GS1280 and ES47 will be announced during H1CY2003.

[2] Newer AlphaStations based on DS Series servers support Fibre Channel Storage.

[3] ATM using an emulated LAN configuration can be used as a cluster interconnect on all AlphaServer systems, except for AlphaServer 300 and 400 systems. ATM is not supported on the DEC Series systems listed or on VAX systems.

[4] Each "Yes" means that this system is supported on this interconnect but cannot be booted as a satellite over this interconnect.

[5] DSSI is not supported on GS, ES, of DS Series AlphaServers.

[6] Each "Yes+Sat" means that this system is supported on this interconnect and can be booted as a satellite node over this interconnect.

[7] Excludes AlphaServer 2000, 2100, 2100A.

[8] Excludes AlphaServer 1000.

[9] AlphaServer 800 only.

[10] Version 7.1 and higher only. Most models provide FDDI booting capability. Refer to system-specific documentation for details.

[11] Using DEFTA only.

[12] Some models may provide slightly different interconnect support. Refer to system-specific documentation for details.

CI Support

OpenVMS Cluster systems can be configured with multiple CI adapters. Table 2 shows the types of adapters that are supported by each system. There can be only one type of adapter configured in a system (with the exception that, with OpenVMS Version 7.1, CIXCD and CIPCA adapters can be configured together in the same system). The maximum number of each type is noted in the table. The CI adapters in a system can connect to the same or different star couplers.

Note: The CIBCA–A adapter cannot coexist with a KFMSA adapter on the same system.

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

Table 2

System - CIxxx	750	780	BCI	BCA-A	BCA-B	XCD	PCA
AlphaServer GS, 8400	–	–	–	–	–	10	10,26[1]
AlphaServer 8200	–	–	–	–	–	–	10,26[1]
AlphaServer ES, 4000, 4100	–	–	–	–	–	–	3[2]
AlphaServer 4000 + I/O expansion	–	–	–	–	–	–	6[3]
AlphaServer DS, 2100A, 1200	–	–	–	–	–	–	3
AlphaServer 2000, 2100	–	–	–	–	–	–	2[4]
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	–

[1]The two numbers represent the support limits for Version 6.2-1H3 and Version 7.1 and higher, respectively.

[2]For three CIPCAs, one must be CIPCA-AA and two must be CIPCA-BA.

[3]Only three can be CIPCA-AA.

[4]Only one can be a CIPCA-BA.

Observe the following guidelines when configuring CIPCA adapters:

- The CIPCA adapter can coexist on a CI bus with CIXCD and CIBCA-B CI adapters and all variants of the HSC/HSJ controller except the HSC50. Other CI adapters cannot be configured on the same CI bus as a CIPCA. HSC40/70 controllers must be configured with a Revision F (or higher) L109 module.
- The CIPCA-AA adapter occupies a single PCI backplane slot and a single EISA backplane slot.
- The CIPCA-BA adapter occupies two PCI backplane slots.

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 systems that can be connected to a star coupler is **16**, regardless of the number of ports.

Memory Channel Support (Version 7.1 and higher only)

Memory Channel is supported on all HP AlphaServer systems starting with the AlphaServer 1000. Observe the following rules when configuring Memory Channel:

- A maximum of eight systems can be connected to a single Memory Channel interconnect.
- Systems configured with Memory Channel adapters require a minimum of 128 megabytes of memory.
- A maximum of two Memory Channel adapters can be configured in a system. Configuring two Memory Channel interconnects can improve the availability and performance of the cluster configuration. Only one Memory Channel adapter may be configured in an AlphaServer 8xxx DWLPA I/O channel configured with any other adapter or bus option. This restriction does not apply to the DWLPB I/O channel, or to DWLPA I/O channels that have no other adapters or bus options.
- Multiple adapters in a system cannot be connected to the same Memory Channel hub.
- Memory Channel adapters *must* all be the same version. Specifically, a Memory Channel V1.5 adapter cannot be mixed with a Memory Channel V2.0 adapter within the same connection.

DSSI Support

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

- Configure the AlphaServer systems shown in Table 1 with KFPSA (PCI to DSSI) adapters. The KFPSA is the highest-performance DSSI adapter and is recommended wherever possible.
- Other supported adapters include:
  - — KFESB (EISA to DSSI) for all AlphaServer systems except 4xxx and 8xxx models
  - — KFESA (EISA to DSSI) for AlphaServer 2100 systems
  - — KFMSB for Alpha XMI systems
  - — KFMSA for VAX XMI systems
  - — KFQSA for VAX Q-bus systems

- KFMSB adapters and KFPSA adapters cannot be configured on the same DSSI bus.
- Up to **24** KFPSAs can be configured on a system.
- Up to **6** KFMSA/Bs can be configured on an XMI bus.
- Up to **12** KFMSA/Bs can be configured in a system.
- Up to four KFESBs can be configured on a system.
- Up to two KFESAs can be configured on a system.
- A mix of one KFESB and one 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.
- Some of the new generation AlphaServer processors will support DSSI. The GS series and the DS20 series will have support. Other DS series and the ES series will not.

Multihost SCSI Storage Support

OpenVMS Cluster Software provides support for multihost SCSI configurations using Alpha systems and SCSI adapters, devices, and controllers. Table 1 shows which systems can be configured on a multihost SCSI bus.

Any HP AlphaStation or AlphaServer system that supports optional KZPSA (fast–wide differential) or KZPBA–CB (ultrawide differential; Version 7.1-1H1 and higher only) adapters can use them to connect to a multihost SCSI bus. Refer to the appropriate system documentation for system specific KZPSA and KZPBA support information. Single-host Ultra SCSI connections with either the KZPBA–CA (ultrawide single-channel adapter) or the KZPBA–CB (ultrawide differential adapter) are supported in Version 6.2–H3 and higher.

Also, any AlphaStation or AlphaServer system except the AlphaServer 4000, 4100, 8200, and 8400 can use embedded NCR–810-based SCSI adapters, or on pre-EV6 hardware platforms the optional KZPAA adapters, to connect to a multihost SCSI bus.

Additionally, DEC 3000 systems can use optional KZTSA (fast–wide differential) adapters to connect to a multihost SCSI bus.

Note: A wide range of SCSI adapters can be used to connect to a single-host SCSI bus. For further information about the complete range of SCSI support, refer to the *OpenVMS Operating System for VAX and Alpha Software Product Description* (SPD 25.01.xx).

HP recommends optional adapters 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 or SCSI–3 disk. Disks must support the following three features:

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

These SCSI disk requirements are fully documented in the *Guidelines for OpenVMS Cluster Configurations* manual. In general, nearly all disk drives available today, from HP or third-party suppliers, support these features. Known exceptions to the range of HP drives are the RZ25 and RZ26F, which do not support tagged command queueing.

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

HSZ series storage controllers can be configured on a mulithost SCSI bus. Refer to the appropriate HSZ storage controller documentation for configuration information. Note that it is not possible to configure tape drives, floppy disks, or CD-ROMs on HSZ controller storage buses when the HSZ is connected to a multihost SCSI bus.

Multihost SCSI buses must adhere to all SCSI–2 or SCSI–3 specifications. Rules regarding cable length and termination must be adhered to carefully. For further information, refer to the SCSI–2 or SCSI–3 specification or the *Guidelines for OpenVMS Cluster Configurations* manual.

Fibre Channel Storage Support

Beginning with Version 7.2–1, OpenVMS Cluster Software provides support for multihost Fibre Channel storage configurations using Alpha systems and Fibre Channel adapters, switches, and controllers. Direct-attached Fibre Channel storage and Arbitrated Loop Fibre Channel configurations are *not* supported. For the current configuration guidelines and limitations, refer to the *Guidelines for OpenVMS Cluster Configurations* manual. This manual outlines the specific requirements for the controller (HSG80, HSG60, and HSV110), switch, and adapter (KGPSA–**), and for the disks that can be attached to this configuration. The number of hosts, adapters, switches, and distance between these

items, is constantly being increased, so refer to the manual for the up-to-date information on this evolving area.

Starting with OpenVMS Version 7.2-2, SCSI tape devices can be connected to a Fibre Channel storage environment with the use of a Modular Data Router (MDR) bridge product. This bridge allows these tape devices to be placed behind the Fibre Channel switch environment and to be shared via the same methodologies as the Fibre Channel disks in the same fabric.

Because the support for Fibre Channel is currently limited to storage only, a second interconnect for node-to-node communications must be present for the full clustered capability to be utilized.

DECamds Console

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

SOFTWARE REQUIREMENTS

OpenVMS Operating System

For information about OpenVMS Operating System Version 7.3, refer to the *OpenVMS Operating System for Alpha and VAX Software Product Description* (SPD 25.01.xx).

For information about OpenVMS Operating System Version 7.2-2, refer to the *OpenVMS Operating System for Alpha and VAX Software Product Description* (SPD 41.87.xx).

The ability to have more than one version of OpenVMS in an OpenVMS Cluster allows upgrades to be performed in a staged fashion so that continuous OpenVMS Cluster system operation is maintained during the upgrade process. Only one version of OpenVMS can exist on any system disk; multiple versions of OpenVMS in an OpenVMS Cluster require multiple system disks. Also, 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 an OpenVMS Cluster configuration is supported according to the following conditions:

- *Warranted support* is provided for mixed-architecture OpenVMS Cluster systems in which all Alpha and VAX systems are running the *same* version of OpenVMS—Version 6.2-xxx, Version 7.0, Version 7.1–xxx, Version 7.2–xxx, or Version 7.3–xxx.

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

- *Migration support* is provided for OpenVMS Cluster systems running *two* versions of the OpenVMS operating system. These versions can be:

  — Any mix of Version 7.3–2, Version 7.3–1, Version 7.3, Version 7.2–2, Version 7.2–1, Version 7.2–1H1, Version 7.2

  — Any mix of Version 7.3–1, Version 7.3, Version 7.2–2, Version 7.2–1xx, Version 7.2, Version 7.1-2, Version 7.1-1Hx, and Version 7.1.

  — Any mix of Version 7.2, Version 7.1–xxx, and Version 6.2–xxx.

  — Any mix of Version 7.1, Version 7.0, and Version 6.2–xxx.

  — Any mix of Version 6.2–xxx with OpenVMS VAX Version 5.5-2, Version 6.0, Version 6.1 and OpenVMS Alpha Version 1.5, Version 6.0, Version 6.1.

  Migration support means that HP 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. HP will answer problem reports submitted about these configurations. However, in exceptional cases, HP may recommend that you move your system to a warranted configuration as part of the solution.

Note: HP does not support the use of more than two versions of OpenVMS software in the same OpenVMS Cluster at the same time. However, in many cases, running more than two versions or mixing versions not described above will operate satisfactorily.

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

DECnet software

DECnet software is not required in an OpenVMS Cluster configuration. However, DECnet software is necessary for internode process-to-process communication that uses DECnet mailboxes.

The OpenVMS Version 6.2–1H3 Monitor utility uses DECnet for intracluster communication.

The OpenVMS Version 7.1 (and higher) Monitor utility uses TCP/IP or DECnet based transports, as appropriate, for intracluster communication.

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

DECamds

DECamds requires HP DECwindows Motif for OpenVMS. For details, refer to the *DECwindows Motif for OpenVMS Software Product Description* (SPD 42.19.xx).

OPTIONAL SOFTWARE

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

Optional products that may be useful in OpenVMS Cluster systems include:

- Volume Shadowing for OpenVMS (SPD 27.29.xx)
- RAID Software for OpenVMS (SPD 46.49.xx)
- DECram for OpenVMS (SPD 34.26.xx)
- VAXcluster Console System (SPD 27.46.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. For more information, refer to the OpenVMS Operating System for VAX and Alpha SPD.

ORDERING INFORMATION

OpenVMS Cluster Software is orderable as follows:

Every server (nonclient) Alpha system in an OpenVMS Cluster configuration requires:

- VMScluster Software for OpenVMS Alpha
  - — Software Licenses: QL–MUZA*–AA
  - — Software Product Services: QT–MUZA*–**
  - — LMF PAK Name: VMSCLUSTER

Note: VMScluster Software for OpenVMS Alpha provides a unique ordering and pricing model for single-CPU and dual-CPU capable systems. Specifically, all AlphaServer DS-series systems, along with AlphaServer 800 and 1200 systems, should use the QL–MUZAC–AA license order number; for service, use the corresponding QT–MUZAC–** order number. For all remaining AlphaServer systems in the Workgroup system class (such as the ES40), use the standard QL–MUZAE–AA license order number; for service, use the corresponding QT–MUZAE–** order number. VMScluster pricing and ordering for the remaining system classes of AlphaServers are unchanged.

Every server (nonclient) VAX system in an OpenVMS Cluster configuration requires:

- VAXcluster Software for OpenVMS VAX
  - — Software Licenses: QL–VBRA*–AA
  - — Software Product Services: QT–VBRA*–**
  - — LMF PAK Name: VAXCLUSTER

OpenVMS Cluster Client Software is available as part of the NAS150 product. It is also separately orderable for DS Series AlphaServers.

- VMScluster Client Software for OpenVMS Alpha
  - — Software Licenses: QL–3MRA*–AA
  - — Software Migration Licenses: QL–6J7A*–AA
  - — Software Product Services: QT–3MRA*–**
  - — LMF PAK Name: VMSCLUSTER-CLIENT

* 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 DECamds and Availability Manager availability management software is included in all the licenses in the preceding list.

DOCUMENTATION

The following manuals are included in the OpenVMS hardcopy documentation as part of the full documentation set:

- *OpenVMS Cluster Systems*
- *Guidelines for OpenVMS Cluster Configurations*
- *DECamds User's Guide*
- *Availability Manager User's Guide*

Refer to the *HP OpenVMS Operating System for VAX and Alpha Software Product Description* (SPD 25.01.xx) for additional information about OpenVMS documentation and how to order it.

Specific terms and conditions regarding documentation on media apply to this product. Refer to HP'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 HP'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 HP of the software to which they relate."

SOFTWARE LICENSING

This software is furnished under the licensing provisions of Hewlett-Packard Company's Standard Terms and Conditions. For more information about HP's licensing terms and policies, contact your local HP 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 *hp 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 HP. For more information, contact your local HP office.

SOFTWARE WARRANTY

This software is provided by HP with a 90 day conformance warranty in accordance with the HP warranty terms applicable to the license purchase.

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

© 2003 Hewlett-Packard Development Company, L.P.

The information contained herein is subject to change without notice. The only warranties for HP products and services are set forth in the express warranty statements accompanying such products and services. Nothing herein should be construed as constituting an additional warranty. HP shall not be liable for technical or editorial errors or omissions contained herein.

Proprietary computer software. Valid license from HP required for possession, use or copying. Consistent with FAR 12.211 and 12.212, Commercial Computer Software, Computer Software Documentation, and Technical Data for Commercial Items are licensed to the U.S. Government under vendor's standard commercial license.