



Software Product Description

PRODUCT NAME: OpenVMS Cluster Software

SPD 29.78.16

This Software Product Description describes Versions 6.2-1H3, 7.1-1H1, 7.1-1H2, 7.1-2, and 7.2 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 all of the above mentioned versions 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 an OpenVMS Cluster.

CPUs 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, CPUs retain their independence because they use local, memory-resident copies of the OpenVMS operating system. Thus, OpenVMS Cluster CPUs can

boot and shut down independently while benefiting from common resources.

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

Because resources are shared, OpenVMS Cluster systems offer higher availability than standalone CPUs. Properly configured OpenVMS Cluster systems can withstand the shutdown or failure of various components. For example, if one CPU in an OpenVMS Cluster 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 OpenVMS Cluster systems have the following software features in common:

- The OpenVMS operating system and OpenVMS Cluster 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 OpenVMS Cluster software using a flexible and sophisticated per-CPU voting mechanism.
- Shared batch and print queues are accessible from any CPU in the OpenVMS Cluster 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 OpenVMS Cluster system.
- All disks and tapes in an OpenVMS Cluster 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 available only with Version 7.1.)
- Configuration command procedures assist in adding and removing CPUs 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 (V7.2 and above).
- An Application Programming Interface (API) that allows applications within multiple OpenVMS Cluster nodes to communicate with each other (V7.2 and above).
- 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 DECamsd availability management tool 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 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 an OpenVMS Cluster node. An OpenVMS Cluster node can be referred to as an OpenVMS Cluster 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 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.
- **HSZ** — An intelligent mass storage controller subsystem that connects to the SCSI bus.
- **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 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 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 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 — CPUs in a OpenVMS Cluster system can 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. An OpenVMS Cluster system must be configured with at least one voting CPU.

OpenVMS Cluster Client Software

OpenVMS Cluster 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 OpenVMS Cluster functionality as described in this SPD, with the following exceptions:

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

Interconnects

OpenVMS Cluster systems are configured by connecting multiple CPUs with a communications medium, referred to as an interconnect. OpenVMS Cluster CPUs 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)
- FDDI (Fiber Distributed Data Interface)
- Ethernet
- Memory Channel (Version 7.1 and above only)

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

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 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, OpenVMS Cluster support for both of these LAN types is identical. Ethernet and FDDI provide CPU-to-CPU communication. Storage can be configured in FDDI environments using FDDI-based storage servers.

OpenVMS Cluster configurations can be configured using wide area networking (WAN) infrastructures, such as DS3, E3, Gigabit Ethernet and ATM. Connection to these media is achieved with FDDI bridges or switches.

Memory Channel is a high-performance interconnect that provides CPU-to-CPU 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 CPUs supported in an OpenVMS Cluster system is **96**.
- Every CPU in an OpenVMS Cluster system must be connected to every other CPU via any of the supported OpenVMS Cluster 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 adapter or through a local CI or DSSI connection. These CPUs cannot be configured to boot as satellite nodes.
- All CPUs 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 CPUs 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 CPUs. 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 CPU-to-CPU communication across the DSSI; CPUs using this adapter must include another interconnect for CPU-to-CPU communication.
- The maximum number of CPUs 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 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 three. If the SCSI bus includes a 5-port (or greater) Fair Arbitration SCSI Hub (i.e. DWZZH-05), the maximum number of CPUs is increased to four.
- The maximum number of multihost SCSI buses that a CPU can be connected to is 12.
- Beginning with OpenVMS V7.2, support for *SCSI Multipath Failover* exists. This feature allows for the failover of cluster storage communications from one path to another when multiple SCSI busses have been connected to the same data source. For detailed information, refer to the *Guidelines for OpenVMS Cluster Configurations* manual.
- OpenVMS Cluster CPUs that are configured using WAN interconnects must adhere to the detailed line specifications described in the *Guidelines for OpenVMS Cluster Configurations* manual. The maximum CPU separation is 150 miles.
- A single time-zone setting must be used by all CPUs 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 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 7.1 *cannot* coexist on a system disk with OpenVMS VAX Version 7.1.
- 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 adapter is not permitted. When two local adapters are used for dual pathing, each adapter 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 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

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.

- The maximum number of bridges between any two end nodes is seven.
- The maximum transit delay through any bridge must not exceed two 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 OpenVMS Cluster traffic on the LAN from any CPU 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 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

Compaq 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 CPUs should be configured using interconnects that provide appropriate performance for the required system usage. In general, use the highest performance interconnect possible. CI and Memory Channel are the preferred interconnects between powerful CPUs.
- 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 CPUs have direct access to shared storage; whenever possible, configure systems to allow direct access to shared storage in favor of OpenVMS MSCP served access. Multi-access CI, DSSI, and SCSI storage provides higher data availability than singly accessed, local adapter-based storage. Additionally, dual pathing of disks between local or HSC/HSJ/HSD/HSZ storage controllers enhances data availability in the event of controller failure.
- OpenVMS Cluster 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 OpenVMS Cluster interconnects (CI, Memory Channel, DSSI, SCSI, Ethernet, and FDDI) 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 multi-path disks will fail over to another path. Disk servers should be the most powerful CPUs 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. 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, nodes may leave the OpenVMS Cluster. The effective performance of the ELAN can be increased by following these guidelines:
 - Configure high-performance CPUs with multiple LAN adapters connected to different LAN segments.
 - Minimize the number of bridges on the path between CPUs that communicate frequently, such as satellites and their boot servers.
 - Use bridges to isolate and localize the traffic between CPUs that communicate with each other frequently. For example, use bridges to separate the OpenVMS Cluster from the rest of the ELAN and to separate CPUs within a cluster that communicate frequently from the rest of the OpenVMS Cluster.

- Use FDDI on the communication paths that have the highest performance requirements. The NISCS_MAX_PKT_SZ 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 high-performance, memory-resident RAM disks. Refer to the DECram for OpenVMS Software Product Description (SPD 34.26.xx) for additional information.

DECamds Features

OpenVMS software incorporates the features of a real-time monitoring, investigation, diagnostic, and system management tool that can be used to improve overall cluster system availability. DECamds can be used in both clustered and non-clustered LAN environments.

The 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.

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 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 for more information.

Interconnect Support

Table 1 shows which CPUs are supported on which interconnects and whether the CPU 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

OpenVMS Cluster systems can use all Ethernet (10Mb/sec and 100Mb/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 only be configured with one FDDI adapter). Refer to the OpenVMS Operating System for VAX and Alpha SPD for more information.

Gigabit Ethernet LAN adapters, supported in V7.1-2 and above, may be used for *limited* OpenVMS Cluster interconnect capability. Refer to the Release Notes for the version being used for specific limitations on this interconnect.

The DEFZA FDDI adapter is supported on VAX systems only.

Note: VAX systems cannot be booted over an FDDI.

Table 1

CPU	CI	Memory Channel ¹	DSSI	SCSI ²	FDDI	Ethernet
AlphaServer GS Series, 8200, 8400	Yes ³	Yes	Yes	Yes	Yes+Sat ⁴	Yes
AlphaServer 4000, 4100	Yes	Yes	Yes	Yes	Yes+Sat	Yes
AlphaServer 1200, 2000, 2100, 2100A	Yes	Yes	Yes	Yes	Yes+Sat	Yes+Sat
AlphaServer 1000, 1000A	-	Yes	Yes	Yes	Yes+Sat	Yes+Sat
AlphaServer 400,800	-	-	Yes	Yes	Yes+Sat ¹	Yes+Sat
AlphaServer 300	-	-	-	Yes	Yes	Yes+Sat
AlphaStations	-	-	-	Yes	Yes+Sat ⁷	Yes+Sat
DEC 7000, 10000	Yes	-	Yes	-	Yes+Sat	Yes
DEC 4000	-	-	Yes	-	Yes	Yes+Sat
DEC 3000	-	-	-	Yes	Yes+Sat ⁵	Yes+Sat
DEC 2000	-	-	-	-	Yes	Yes+Sat
VAX 6000, 7000, 10000	Yes	-	Yes	-	Yes	Yes
VAX 8xxx, 9xxx, 11/xxx	Yes	-	-	-	-	Yes
VAX 4xxx ⁶	-	-	Yes	-	Yes	Yes+Sat
VAX 2xxx, 3xxx ⁶	-	-	-	-	-	Yes+Sat

¹Version 7.1 and above only.

²This column refers to multihost SCSI connectivity. Refer to the appropriate system documentation for information regarding single-host connectivity to SCSI buses.

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

⁴Each "Yes+Sat" means that this CPU is supported on this interconnect and can be booted as a satellite node over this interconnect.

⁵Using DEFTA only.

⁶Some models may provide slightly different interconnect support. Refer to system-specific documentation for details.

⁷Version 7.1 and above only. Most models provide FDDI booting capability. Refer to system-specific documentation for details.

CI Support

OpenVMS Cluster CPUs 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 in a CPU (with the exception that, with V7.1, CIXCD and CIPCA adapters can be configured together in the same 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 system.

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

Table 2

CPU - Cxxxx	750	780	BCI	BCA-A	BCA-B	XCD	PCA
AlphaServer 8400	-	-	-	-	-	10	10,26 ¹
AlphaServer 8200	-	-	-	-	-	-	10,26 ¹
AlphaServer 4000, 4100	-	-	-	-	-	-	3 ²
AlphaServer 4000 + I/O expansion	-	-	-	-	-	-	6 ³
AlphaServer 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	-

¹The two numbers represent the support limits for Version 6.2-1H3 and Version 7.1 (and above) respectively.

²For 3 CIPCA's, one must be a CIPCA-AA and two must be CIPCA-BA

³Only 3 may be CIPCA-AA

⁴Only one may 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 CPUs that can be connected to a star coupler is **16**, regardless of the number of ports.

Memory Channel Support—Version 7.1 and above Only

Memory Channel is supported on all AlphaServer systems from (and including) the AlphaServer 1000 upwards. Observe the following rules when configuring Memory Channel:

- A maximum of eight CPUs can be connected to a single Memory Channel interconnect.
- CPUs 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 CPU. 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 CPU cannot be connected to the same Memory Channel hub.
- Memory Channel adapters MUST be all of the same version level. Specifically, a V1.5 Memory Channel 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 below). 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 six 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.

Multihost SCSI 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 AlphaStation or AlphaServer system that supports optional KZPSA (fast-wide differential) or KZPBA-CB (Ultra-wide differential - V7.1-1H1 and above 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 (Ultra-wide single ended adapter) or the KZPBA-CB (Ultra-wide differential adapter) are supported in V6.2-H3 and above.

Also, any AlphaStation or AlphaServer system except the AlphaServer 4000, 4100, 8200, and 8400 can use embedded NCR-810-based SCSI adapters or 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 SPD.

Compaq 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 the *Guidelines for OpenVMS Cluster Configurations* manual. In general nearly all disk drives available today, from Compaq or third party suppliers, support these features. Known exceptions to the range of Compaq 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 multihost 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. Refer to the SCSI-2 or SCSI-3 specification or the *Guidelines for OpenVMS Cluster Configurations* manual for further information.

DECamds Console

Compaq 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

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 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 or Version 7.1.

Warranted support means that Compaq 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:

1. Any mix of Version 7.2, Version 7.1-2, Version 7.1-1Hx, Version 7.1, Version 7.0, and Version 6.2-xxx.
2. 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 Compaq 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. Compaq will answer problem reports submitted about these configurations. However, in exceptional cases, Compaq may recommend that you move your system to a warranted configuration as part of the solution.

Note: Compaq 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.

Compaq 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 above) Monitor utility uses TCP/IP or DECnet based transports, as appropriate, for intracluster communication.

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

DECamds Availability Manager

The DECamds Availability Manager requires DECwindows Motif for OpenVMS (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)
- 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 an OpenVMS Cluster configuration requires:

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

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 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 DECamds Availability Manager is included in all the licenses in the preceding list.

DOCUMENTATION

The *OpenVMS Cluster Systems* manual, the *Guidelines for OpenVMS Cluster 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 Compaq'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 Compaq'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 Compaq of the software to which they relate."

SOFTWARE LICENSING

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

YEAR 2000 READINESS

This product is capable of accurately processing, providing, and/or receiving date data from, into and between the twentieth and the twenty-first centuries, and the years 1999 and 2000, including leap year calculations, when used in accordance with the associated product documentation and provided that all hardware, firmware and software used in combination with the product properly exchange accurate date data with the product.

SOFTWARE PRODUCT SERVICES

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

SOFTWARE WARRANTY

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

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

© 1999 Compaq Computer Corporation.
All rights reserved.

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

NCR is a registered trademark of NCR Corporation.

