# Software Product Description

**PRODUCT NAME: VSI OpenVMS Cluster Software** **DO-VIBHAA-032**

This Software Product Description describes Versions 8.4-1H1, 8.4-2, and 8.4-2L1 of the VSI OpenVMS Cluster Software for Integrity servers.

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 Integrity server systems, or in mix of AlphaServer and Integrity server systems. In this SPD, this environment is referred to as an OpenVMS Cluster system.

*Note: VSI OpenVMS versions do not support clusters that contain a VAX system.*

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.
- 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.
- An application programming interface (API) allows applications within multiple OpenVMS Cluster nodes to communicate with each other.
- 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.
- OpenVMS Clusters Software features the Cluster over IP functionality, which provides the ability to form clusters beyond a single LAN or VLAN segment using industry standard Internet Protocol. This feature provides improved disaster tolerant capability.
- The optionally installable Availability Manager (as well as the DECamds availability management tool) allows system managers to monitor and manage resource availability in real time on all the members of an OpenVMS Cluster.
- Satellite boot support on Integrity server systems is available.
- 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 Maintenance Operations Protocol (MOP) server and a disk server. A boot node can fully service satellite boot requests.
- System — An Integrity server or an AlphaServer 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 — Intelligent mass storage controller sub- systems that connect to the CI bus.
- HSD — An intelligent mass storage controller sub- system that connects to the DSSI bus.
- HSG, HSV/EVA, MSA, XP — Intelligent mass storage controller subsystems that connect to the Fibre Channel bus.
- HSZ — An intelligent mass storage controller sub- system 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 pro- vide 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 an AlphaServer and Integrity server 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.
- OpenVMS Cluster over IP and IP Cluster Interconnect (IPCI) terms are interchangeably and refers to using TCP/IP stack for cluster communication.
- 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. 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:

- Fibre Channel (storage only, Version 7.2–1 and higher only, Integrity servers and Alpha)
- Ethernet (10/100, Gigabit,10 Gigabit) (Integrity servers and Alpha)
- SCSI/SAS (Small Computer Storage Interconnect/Serial attached SCSI) (Alpha (SCSI only) and Integrity servers)
- Cluster over IP (IPCI) (Alpha and Integrity servers)
- Virtual Connect (Integrity servers only)
- CI (computer interconnect) (Integrity servers and Alpha)
- DSSI (Digital Storage Systems Interconnect) (Integrity servers and Alpha)
- FDDI (Fiber Distributed Data Interface) (Integrity servers and Alpha)

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 industry standard interconnect for storage and communications and allows for a storage-only interconnect in a multihost environment utilizing Fibre Channel switched topologies. SCSI tapes utilizing the Modular Data Router bridge or the Network Storage Router bridge are 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.

Cluster over IP provides the ability to form clusters beyond a single LAN or VLAN segment using industry standard Internet Protocol. This feature provides improved disaster tolerant capability. System managers also have the ability to manage or monitor OpenVMS cluster that uses IP for cluster communication using SCACP management utility.

Cluster protocol (SCS also known as SCA) over LAN is provided by Port Emulator driver (PEDRIVER). PEDRIVER uses User Datagram Protocol (UDP) and IP in addition to directly using 802.3 interfacing with LAN for cluster communication. The datagram characteristics of UDP combined with PEDRIVER's inbuilt reliable delivery mechanism are used for transporting cluster messages that are used by SYSAP (system level application) to communicate between two cluster nodes.

Cluster over IP is an optional feature that can be enabled in addition to the traditional LAN based communication. However, if both LAN and IP mode of communication exist between nodes in a cluster, PEDRIVER prefers LAN communication instead of IP.

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*

- Mixed-architecture clusters are limited to a mixed AlphaServer and Integrity server cluster.
- The maximum number of systems supported in an OpenVMS Cluster system is **96**. With OpenVMS Version 8.2 for Integrity servers, a restriction exists where the maximum mixed-architecture cluster node count is **16**. For OpenVMS V8.2-1 for Integrity server systems and later, this restriction has been removed.
- Every system in an OpenVMS Cluster system must be connected to every other system via any supported OpenVMS Cluster interconnect.
- 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. 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 configuration manuals for further information.
- The maximum number of AlphaServer 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 AlphaServer systems is increased to **4**. A maximum of **2** Integrity server systems can be connected to a SCSI bus. This

configuration has several restrictions, which are documented in *HPE OpenVMS Version 8.2–1 for Integrity Servers New Features and Release Notes*.

- The maximum number of multihost SCSI buses to which an AlphaServer system can be connected to is **26**.
- The configuration size for Fibre Channel storage increases on a regular basis with new updates to OpenVMS. Please refer to the *Guidelines for OpenVMS Cluster Configurations* manual for the most up-to-date configuration capabilities.
- Multipath Failover for both Parallel SCSI and Fibre Channel storage environments is supported. This feature allows failover from a locally connected storage path to a served path for data access. For detailed information, refer to the *Guidelines for OpenVMS Cluster Configurations* manual.
- Multipath failover to the MSCP served path is also 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. With proper consulting support from VSI Professional Services, the maximum system separation is 500 miles.
- Using VSI TCP/IP Services for OpenVMS Version 5.7, OpenVMS cluster systems can use IP for cluster communication.
- 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 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 8.4 *cannot* coexist on a system disk with OpenVMS Version 8.4-2H1.
- 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 HSZcontrollers 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 for Integrity servers software installations and upgrades must be performed using an Integrity server 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**.

[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 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

VSI 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. Virtual Connect is available on HPE BladeSystem c-Class enclosures.

- 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. Multi-access 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 or HSC/HSJ/HSD/HSZ/HSG/MSA/XP/EVA 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 band- width 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:
  — *VSI 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 shadow- set members

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

- The 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 DO-VIBHAB-005) 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.

*Availability Manager*

Availability Manager is a system management tool that enables one or more OpenVMS Alpha nodes to be monitored on an extended local area network (LAN) from an OpenVMS Alpha, OpenVMS for Integrity servers, or a Windows node. This tool helps system managers and analysts target a specific node or process for detailed analysis. The analysis detects resource availability problems and suggests corrective actions.

Availability Manager has been enhanced with the ability to support the OpenVMS Cluster over IP functionality. Availability Manager provides the functionality to manage and monitor LAN or IP path (channels) data.

*SCACP*

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

**HARDWARE SUPPORT**

*System Support*

VSI supports VSI OpenVMS clusters containing Integrity servers and AlphaServer systems.

*Peripheral Option and Storage Controller Support*

OpenVMS Cluster systems can use all peripheral options and storage subsystems supported by OpenVMS. Refer to the VSI OpenVMS Operating System Software Product Description (SPD DO-VIBHA*-005) 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, 100 Mb/sec, and 1000 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 VSI OpenVMS Operating System Software Product Description (SPD DO-VIBHA*-005) for more information.

**AlphaServer and Integrity Server Cluster Interconnect Support**

Starting with the OpenVMS V8.2 release of OpenVMS Cluster software, support was provided for mixed-architecture Alpha and Integrity Server clusters. A limited number of interconnects exist for this environment.

For Host-to-host (SCS) communications, AlphaServer and Integrity servers can be connected via 10/100/1000 LAN-based connections. For shared Storage, the primary environment is Fibre-Channel SAN-based fabric.

A 2-node shared SCSI environment is provided via the MSA30MI shelf between Integrity Server nodes ONLY. On Integrity servers, support for up to 4-node shared SAS cluster environments is provided via the MSA60/70 shelf and connected by the Smart Array P700 controller.

*CI Support*

OpenVMS Cluster systems can be configured with multiple CI adapters. The table below shows the types of adapters that are supported by each system. There can be only one type of adapter configured in a system. However, 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 1: Types of Adapters Supported by System**

| System - CIxxx | | | | | | | |
|---|---|---|---|---|---|---|---|
| **AlphaServer GS, 8400** | – | – | – | – | – | 10 | 10,26[1] |
| **AlphaServer 8200** | – | – | – | – | – | – | 10,26[1] |
| **AlphaServer ES, 4000, 4100** | – | – | – | – | – | – | 3[2] |
| **AlphaServer 4000 + I/O** | – | – | – | – | – | – | 6[3] |
| **AlphaServer DS, 2100A, 1200** | – | – | – | – | – | – | 3 |
| **AlphaServer 2000, 2100** | – | – | – | – | – | – | 2[4] |
| **DEC 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 back- plane 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*

Memory Channel is supported on all HPE 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 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
- 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.

*Dual-Host SCSI Storage Support (for Integrity servers)*

OpenVMS Integrity server systems that are members of the rx1600 and rx2600 families can form a 2-node OpenVMS for Integrity servers Cluster system with shared SCSI storage. The configuration has many restrictions as documented in *HPE OpenVMS Version 8.2–1 for Integrity Servers New Features and Release Notes*.

*Multihost SCSI Storage Support (Alpha)*

OpenVMS Cluster Software provides support for multihost SCSI configurations using AlphaServer systems and SCSI adapters, devices, and controllers.

Any HPE 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 *VSI OpenVMS Operating System for 8.4-2 Software Product Description* (SPD DO-VIBHAB-005).

HPE 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, many disk drives available today, from HPE or third-party suppliers, support these features. Known exceptions to the range of HPE 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*

OpenVMS Cluster Software provides support for multihost Fibre Channel storage configurations using AlphaServer systems and Fibre Channel adapters, switches, and controllers. This support is also available for Integrity servers. 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, HSV110, MSA1000 and XP), 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.

SCSI tape devices can be connected to a Fibre Channel storage environment with the use of a Modular Data Router (MDR) or Network Storage Router (NSR) bridge products. These bridges allow 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. This support is also available for Integrity servers.

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.

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.

*Virtual Connect Support*

Virtual Connect is a set of interconnect modules and embedded software for HPE BladeSystem c-Class enclosures; it simplifies the setup and administration of server connections. HPE Virtual Connect includes the HPE 1/10Gb Virtual Connect Ethernet Module for c-Class BladeSystem, the HPE 4Gb Fibre Channel module, and the HPE Virtual Connect Manager.

**Warranted and Migration Cluster Support**

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

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

**Note:** Same platform/version pairings are always warranted.

## SOFTWARE REQUIREMENTS

*OpenVMS Operating System*

*Migration support* is also provided for OpenVMS Cluster systems running *two* versions of the OpenVMS operating system

*DECnet software (Alpha and Integrity servers)*

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

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

*DECamds (Alpha and Integrity Servers)*

DECamds requires DECwindows Motif for OpenVMS. For details, refer to the *DECwindows Motif for OpenVMS Software Product Description* (DO-VIBHAA-026).

## 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 DO-VIBHAA-031)
- DECram for OpenVMS (DO-VIBHAB-005)

## 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 latest VSI OpenVMS Operating System SPDs.

## ORDERING INFORMATION

Contact sales@vmssoftware.com for the latest ordering information for VSI OpenVMS.

## DOCUMENTATION

Please refer to the following manuals for more information.

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

Refer to the *VSI OpenVMS for Integrity Servers* Software Product Description 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 VSI 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 VSI copyright notice. Customization or modifications of any kind to the software documentation files are not permitted.

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

## SOFTWARE LICENSING

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

*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 *VSI OpenVMS Version 8.4-2 for Integrity servers Software Product Description* (DO-VIBHAB-005) or the VSI OpenVMS License Management Utility Manual (DO-VIBHAA-004).

## SOFTWARE PRODUCT SERVICES

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

http://vmssoftware.com/services

## SOFTWARE WARRANTY

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