

Software Product Description

PRODUCT NAME: VAXcluster Software, Version 5.4

SPD 29.78.04

DESCRIPTION

VAXcluster Software is a VMS System Integrated Product (SIP). It provides a highly integrated VMS computing environment distributed over multiple VAX and/or MicroVAX CPUs. This environment is called a VAXcluster system.

CPUs in a VAXcluster system can share processing, mass storage, and other resources under a single VMS security and management domain. Within this highly integrated environment, CPUs retain their independence because they use local, memory-resident copies of the VMS Operating System. Thus, clustered CPUs can boot and fail independently while benefiting from common resources.

Applications running on one or more CPUs in a VAXcluster system access shared resources in a coordinated manner. VAXcluster Software components synchronize access to shared resources, preventing two or more processes from interfering with each other when updating data. This coordination ensures data integrity during multiple concurrent update transactions.

Because resources are shared, VAXcluster systems offer higher availability than standalone CPUs. Properly configured VAXcluster systems can withstand the failure of various components. For example, when one CPU in a cluster fails, users can log on to another CPU to create a new process and continue working. Applications can be designed to survive the failure of one or more CPUs or other components.

All VAXcluster systems have the following software features in common:

- Shared mass storage - the VMS file system allows all CPUs in a VAXcluster system to share mass storage, whether the storage devices are connected to a Hierarchical Storage Controller (HSC) subsystem or to a CPU.

- Shared batch and print queues accessible to any CPU in the VAXcluster system. The VMS job controller manages 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 that the batch load is shared.
- VMS MSCP server load balancing - the first time the disk is mounted it is brought online to the least loaded server.
- All disks in the VAXcluster system can be made accessible to all CPUs.
- Process information and control are available clusterwide to application programs and system utilities.
- The automated configuration command procedure assists in adding and removing CPUs and in modifying their configuration characteristics.
- Dynamic Show Cluster Utility displays the status of cluster hardware components and cluster communication links.
- There are cluster-wide standard VMS system and security features.

VAXcluster System Configurations

VAXcluster Software supports combinations of the following interconnects:

- Computer Interconnect (CI)
- Ethernet
- Digital Small Storage Interconnect (DSSI)

The optimal VAXcluster system for any computing environment is based on tradeoffs of cost, functionality, and performance. Performance is influenced by the following factors:

- Applications in use
- Number and model of CPUs
- Interconnects and adapter types
- Disk I/O capacity and access time
- Number of disks being served

- Interconnect utilization

Definitions:

In VAXcluster systems, some CPUs can perform specialized functions. The following terms are used to refer to those CPUs:

- A VAX CPU (Central Processing Unit) - This is an electronic unit that includes one or more processors and acts as a central controlling body.
- Disk server - A disk server is a VAX system that makes disks available to other systems in the VAXcluster system.
- Boot server - A boot server is a disk server that down-line loads the VMS Operating System to other cluster members over the Ethernet.
- Satellite - A satellite is a system that is booted over the Ethernet by a boot server. The satellite systems are listed in the *HARDWARE REQUIREMENTS* section.
- Local Area VAXcluster System - This is an Ethernet-based VAXcluster system.
- Mixed Interconnect VAXcluster System - This type of VAXcluster system uses more than one type of interconnect for cluster communication.

The following general configuration rules apply to VAXcluster systems:

- The maximum number of VAX CPUs supported in a VAXcluster system is 96. Up to 32 systems may be systems other than single user workstations.
 - If the Ethernet is used for cluster communications in any VAXcluster system, all CPUs must be connected to the same Ethernet local area network. VAXcluster systems support one Ethernet adaptor per CPU; all cluster members must be reachable on the Ethernet via that adapter. The adaptor used by the VAXcluster software is not user selectable. CPUs may have other adapters for other uses.
 - A star coupler is a common connection point for CI connected CPUs and HSC subsystems. A VAXcluster system may include multiple star couplers.
 - The RA series disks may be dual pathed between pairs of HSC subsystems on the same star coupler or between a pair of local controllers. Such dual-pathing provides both enhanced data availability and failover capability. Failover occurs when one controller or cable malfunctions and breaks one path. When the path breaks, the device using that path automatically fails over to the other path. The failover process is transparent to applications. There cannot be dual porting between an HSC and a local adaptor.
- In Dual-Host DSSI configurations, both CPUs must be configured as VAXcluster members, and the Digital Storage Systems Interconnect (DSSI) must be connected to both CPUs using any combination of embedded adaptors or KFQSA adaptors. Both members of the dual host system must be connected to the same extended Local Area Network (LAN).
 - CPUs may use the Ethernet for cluster communications and may concurrently use it for other network protocols that conform to the applicable Ethernet standards, such as Ethernet V2.0, IEEE 802.2 and IEEE 802.3.
 - Ethernet bridges may be used to localize VAXcluster system traffic should the overall network traffic be of concern. It is possible for a VAXcluster system to exist on both sides of the bridge. However, a low-latency data path providing approximately 10 megabits per second throughput must link all clustered CPUs.
 - A CI path is needed between any pair of the following CPUs in a VAXcluster system. An Ethernet only VAXcluster system may include a maximum of one of the following CPUs:
 - VAX 6000-xxx
 - VAX 85xx
 - VAX 86xx
 - VAX 8700
 - VAX 88xx
 - VAX 9xx

Recommendations

Digital recommends VAXcluster system configurations based on its experience with the VAXcluster Software Product. The customer should evaluate specific application dependencies and performance requirements to determine an appropriate configuration for the desired computing environment. The requirements and recommendations are applicable at the time of release, but may change over time. When planning a VAXcluster system, consider the following recommendations:

- While VAXcluster systems may include any number of system disks, performance and disk space should be considered in determining their number and location. It is important to recognize that system management efforts increase in proportion to the number of system disks.
- VAXcluster systems can provide enhanced availability by utilizing redundant components.

- If possible, Local Area and mixed-interconnect VAXcluster systems should include multiple boot and disk servers to enhance availability. When a server fails in configurations that include multiple servers, satellite access to disks fails over to another server. Boot and disk servers should be the most powerful CPUs in the cluster and should use the highest-bandwidth Ethernet adapters available in the cluster.
- Dual-pathing of DSA disks between HSC subsystems or between local storage adapters enhances availability. In Local Area VAXcluster systems, DSA disks can be dual pathed between local storage adapters. RF-series disks can be accessed by both CPUs in Dual-Host configurations.
- The optional VMS Volume Shadowing System Integrated Product provides the following advantages:
 - Increased data availability
 - Enhanced read performance with multiple system and data disks

For more information, refer to the VMS Volume Shadowing Software Product Description (SPD 27.29.xx.)
- A dual-host configuration is a pair of MicroVAX CPUs connected to the same Digital System Storage interconnect (DSSI) bus. For the MicroVAX 3500 and MicroVAX 3600, the console ROMS must be Revision Level V5.1
- Dual-Host configurations utilizing RF-series Integrated Storage Elements (ISEs) offer failover and increased availability:
 - CPUs have concurrent access to any disk on the DSSI bus.
 - An ISE single disk is accessed through two paths and can be served to all satellites by either CPU.
 - If either CPU fails, satellites can access their disks through the remaining one.

HARDWARE REQUIREMENTS

Table 1

Processors Supported	Satellite Yes/No
VAXft 3000-xxx	Yes
VAX 4000-xxx	Yes
VAX 6000 Model 200 Series *	No
VAX 6000 Model 300 Series *	No

Processors Supported	Satellite Yes/No
VAX 6000 Model 400 Series	No
VAX 8200	No
VAX 8250	No
VAX 8300	No
VAX 8350	No
VAX 8500	No
VAX 8530	No
VAX 8550	No
VAX 8600	No
VAX 8650	No
VAX 8700	No
VAX 8800	No
VAX 8810	No
VAX 8820	No
VAX 8830	No
VAX 8840	No
VAX 9000-210	No
VAX 9000-410	No
VAX-11/750	No
VAX-11/780	No
VAX-11/785	No
MicroVAX II	Yes
MicroVAX 2000	Yes
MicroVAX 3100	Yes
MicroVAX 3300	Yes
MicroVAX 3400	Yes
MicroVAX 3500	Yes
MicroVAX 3600	Yes
MicroVAX 3800	Yes
MicroVAX 3900	Yes
VAXstation II	Yes
VAXstation 2000	Yes
VAXstation 3100	Yes
VAXstation 3200	Yes
VAXstation 3500	Yes
VAXstation 3520	Yes
VAXstation 3540	Yes
VAXserver 3100	Yes
VAXserver 3300	Yes

Table 1 (Cont.)

Processors Supported	Satellite Yes/No
VAXserver 3400	Yes
VAXserver 3500	Yes
VAXserver 3600	Yes
VAXserver 3602	Yes
VAXserver 3800	Yes
VAXserver 3900	Yes
VAXserver 4000	Yes
VAXserver 6000-210	No
VAXserver 6000-220	No
VAXserver 6000-310	No
VAXserver 6000-320	No
VAXserver 6000-410	No
VAXserver 6000-420	No

* The VAX 6000 Model 200 Series and VAX 6000 Model 300 Series were formerly called the 62xx and 63xx series.

Quantity CI Interfaces Supported by CPUs

Table 2 shows the types of interfaces that can be supported by each CPU. There can only be one type of interface on a system and the maximum quantity of that type is indicated.

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

Table 2

CPU Type	CI750	CI780	CIBI	CIBCA-A	CIBCA-B	CIXCD
11/750	1	-	-	-	-	-
11/780	-	1	-	-	-	-
11/785	-	-	-	-	-	-
62xx, 63xx	-	-	-	1	4	4
6000-2xx	-	-	-	-	-	-
6000-3xx	-	-	-	-	-	-
6000-4xx	-	-	-	-	-	-
82xx 83xx	-	-	1	1	1	-
85xx 8700 88xx	-	-	1	1	2	-
86xx	-	2	-	-	-	-
9xx	-	-	-	-	-	4

CPUs Not Supported:

VAX: VAX-11/725, VAX-11/730, VAX-11/782

MicroVAX: MicroVAX I

VAXstation: VAXstation I, VAXstation 8000

Supported Preconfigured VAXcluster Systems:

All

Supported Ethernet Adapters:

- DEBNA
- DEBNI
- DELQA
- DELUA
- DEMNA
- DEQNA ¹
- DESQA
- DESPA
- DEUNA

Disk and Memory Requirements:

The smallest supported boot server system disks are RD54s, RF30s and RZ23s.

All VAXcluster system CPUs must have a minimum of 4 megabytes of physical memory.

¹ All systems utilizing a DEQNA must operate with software data checking enabled. Since AUTOGEN automatically sets the correct parameter, no system management intervention is required. The DEQNA is not supported in clusters of greater than 42 nodes. For improved performance, Digital recommends replacing DEQNA's with DELQA's. An upgrade program is currently available from Digital. For more information, contact your local sales office.

OPTIONAL HARDWARE

This section describes the Computer Interconnect Star Coupler Expander (CISCE) and Hierarchical Storage Controller (HSC) Subsystems.

A Computer Interconnect Star Coupler Expander (CISCE) may be added to any Star Coupler in a CI-based or mixed-interconnect VAXcluster system to increase the total combined number of CI-connected CPUs and HSC subsystems to 32.

HSC Subsystems

VAXcluster systems may include HSC40, HSC50 and HSC70 intelligent mass storage controllers. These controllers offload disk management functions from host CPUs and provide several additional benefits:

- Increased I/O throughput
- Multiple-CPU access to storage

Support for VMS Volume Shadowing Software

The following rules apply for HSC subsystems:

- Each HSC40 supports a maximum of 3 channels for disks and/or tapes.
- Each HSC50 supports a maximum of 6 channels for disks and/or tapes.
- Each HSC70 supports a maximum of 8 channels for disks and/or tapes.
- Each disk channel supports a maximum of 4 disk drives and each tape channel supports a maximum of 4 tape subsystems (with 4 tape drives each). The VMS Operating System supports a maximum of 12 tape drives per HSC subsystem.
- TA-series tape drives may only be dual pathed between pairs of HSC subsystems with HSC microcode revision level V3.9 or later. Tape drives may not be dual pathed between local adaptors.
- If disks or tapes are dual pathed between a pair of HSC subsystems, both HSC subsystems must be connected to the same Star Coupler.

SOFTWARE REQUIREMENTS

- VMS Operating System

V5.3 and V5.4 of the VMS Operating System may coexist in a VAXcluster with the replacement of an image on the V5.3 system with one supplied with the V5.4 system. See the V5.4 release notes for details.

V5.3 and V5.4 of the VMS Operating System may coexist in a VAXcluster. However, only one version of VMS may exist on a specific system disk. During a rolling upgrade, a separate system disk is required for each version. However, Digital recommends that all VAX systems in a VAXcluster run the latest release of VMS.

- DECnet-VAX Software

All VAXcluster CPUs require either an End Node or Full Function DECnet-VAX license.

Refer to the DECnet-VAX Software Product Description document (SPD 25.03.xx) for further information.

OPTIONAL SOFTWARE

For information on VAXcluster support for optional software products, refer to the *CLUSTER SUPPORT* section of the Software Product Description (SPD) documents for those products.

Optional products that are particularly useful in VAXcluster systems include:

VMS Volume Shadowing Software (SPD 27.29.xx)
 VAX Performance Advisor (SPD 27.71.xx)
 VAXcluster Console System (SPD 27.46.xx)

ORDERING INFORMATION

Software Licenses: QL-VBRA*-AA
 Software Product Services: QT-VBRA*-**

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

SOFTWARE LICENSING

A VAXcluster Software license is required for each CPU in a VAXcluster system.

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

LICENSE MANAGEMENT FACILITY SUPPORT

The VAXcluster Software product supports the VMS License Management Facility.

For more information about the License Management Facility, refer to the VMS Operating System Software Product Description (SPD 25.01.xx) or the License Management Facility Manual in the VMS Operating System documentation set.

For more information about Digital's licensing terms and policies, contact your local Digital office.

SOFTWARE PRODUCT SERVICES

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

SOFTWARE WARRANTY

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

™ The DIGITAL Logo, DECnet-VAX, DELUA, DEQNA, DEUNA, HSC, MicroVAX, RD54, VAX, VAXserver, VAXstation, VAXcluster and VMS are trademarks of Digital Equipment Corporation.