

Software Product Description

PRODUCT NAME: VAXcluster Software, Version 5.1

SPD 29.78.01

DESCRIPTION

VAXcluster Software provides a highly integrated VMS computing environment distributed over multiple VAX and/or MicroVAX systems.

VAXcluster members can share many resources, such as disk and tape storage, CPU resources, and system management operations. Within this highly integrated environment, however, systems retain their independence because they use local, memory-resident copies of VMS. Thus, members can boot and fail independently while benefiting from common resources.

Applications running on the same or multiple nodes in a VAXcluster system access shared resources in a coordinated manner. VAXcluster Software components make it possible to synchronize access to shared resources, preventing two or more processes from interfering with each other when attempting simultaneous reads and/or updates of data. This coordination ensures data integrity during multiple, concurrent update transactions.

Because resources are shared, VAXcluster systems offer higher availability than standalone processors. When one system in a cluster fails, users can log on to another cluster node to create a new process and continue working. Applications can be designed to continue running despite the failure of a VAXcluster node.

VAXcluster Software operates in the following VAXcluster configurations:

- CI-based VAXcluster systems, utilizing the Computer Interconnect (CI) for cluster communication
- Ethernet-based (Local Area VAXcluster) systems, utilizing the Ethernet for cluster communication
- Mixed-interconnect VAXcluster systems, utilizing both the Ethernet and CI for cluster communication

All configurations have the following software features in common:

- Shared disk storage. The VMS file system allows all VAX processors in a VAXcluster system to share disk mass storage, whether the disk is connected to an HSC-series controller or to a VAX processor.

- Clusterwide batch and print queues, accessible to any node in the configuration. The VMS job controller allows clusterwide batch and print queues, which can be accessed by any node within the VAXcluster configuration. Jobs submitted to clusterwide queues are routed to the least loaded system, balancing the batch load across the cluster members.
- All VAXcluster members may have access to all disks in the cluster.
- Automated configuration command procedure to assist in adding, removing, and modifying nodes in the cluster.
- Dynamic Show Cluster utility displaying the status of cluster members, including the status of each node, cables, connections, SCS traffic, and Connection Manager traffic.
- Standard VMS system and security features.

Configurations

DIGITAL recommends configurations based on its experience with the VAXcluster Software product. Specific application dependencies and performance requirements should also be evaluated to determine an appropriate configuration for the desired computing environment.

Some VAXcluster members perform specialized functions. The following terms are used to refer to those nodes in local area and mixed-interconnect configurations. Members not defined in any of these categories are simply equal participating members.

- Satellite nodes are Ethernet-only systems without their own local VMS system disk. They are down-line loaded via the Ethernet by boot servers. Only the processors listed in the MicroVAX and VAXstation categories in the Hardware Requirements section may operate as satellite members.
- Disk servers are VAX systems that make disks available to cluster members. Disks can be served either by HSCs or by VAX nodes running MSCP server software.
- Boot servers are disk servers that down-line load satellites.

digital
software

January 1989

AE-LS19B-TE

The following rules apply to Ethernet-based VAXcluster configurations:

- The cluster may include a maximum of 42 CPUs.
- CPUs use the Ethernet for cluster communications and 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.
- A data path providing approximately 10 megabits per second throughput is recommended between all nodes in a cluster.
- The cluster may include a maximum of one of the following CPUs:
 - VAX 62xx
 - VAX 85xx
 - VAX 86xx
 - VAX 87xx
 - VAX 88xx
- The cluster may include a maximum of two of the following CPUs:
 - VAX 82xx
 - VAX 83xx
- The two previous categories are mutually exclusive - that is, a cluster may not contain one 8600 and one 8250.
- A VAXstation 3500 may act as a boot server for up to five satellite members.
- A VAX-11/750 may act as a boot server for up to 13 satellite members, including a maximum of three MicroVAX 3xxx satellites.
- Each RF30 system disk in a cluster can support up to 5 satellites.
- An RD54-based processor may act as a boot server for up to five satellite members. The following are exceptions:
 - A VAXstation 2000 may act as a boot server for up to three other VAXstation 2000 satellite members.
 - A VAXstation 8000 may act as a boot server for up to three VAXstation II, MicroVAX II or MicroVAX 2000 satellite members.
 - A MicroVAX 2000 may act as a boot server for up to five VAXstation or MicroVAX 2000 satellite members.
 - A VAXstation II may act as a boot server for up to five satellite members, none of which may be a MicroVAX 3xxx system.

The following rules apply to CI-based and mixed-interconnect configurations:

- There may be a maximum of 42 nodes in a mixed-interconnect configuration, with a maximum of 16 CI-connected VAX nodes.
- CI-based configurations may contain a maximum of 16 nodes, including VAX nodes and HSCs.
- A Computer Interconnect Star Coupler Expander (CISCE) may be added to a CI-based or mixed-interconnect VAXcluster system to increase the total number of CI-connected nodes to 24. (See restrictions in Optional Hardware section.)
- There may be a maximum of one Star Coupler, and all CI-connected VAX systems must be connected to it.
- All VAX 8xxx-series and VAX 62xx-series CPUs participating in mixed-interconnect configurations must connect to the CI.
- If the Ethernet is used for cluster communications between any VAXcluster members, all VAXcluster members must connect to it.

When planning a cluster configuration, consider the following recommendations:

- Use systems with the maximum power as boot and/or disk servers.
- While there may be any number of system disks in a cluster, performance and disk space should be considered in determining the number of system disks. It is important to remember that system management costs increase in proportion to the number of system disks.
- Applications that demand high availability benefit most from configurations utilizing the Computer Interconnect (in either a CI-based or mixed-interconnect configuration). Dual-pathing of disks between HSCs or between local storage adapters also enhances availability. Ethernet-based configurations may take advantage of data disks dual pathed between local storage adapters.
- Dual host configurations of the MicroVAX 3300 or MicroVAX 3400 (with both CPUs configured as boot servers) that utilize the RF-series disk drives offer failover and increased availability:
 - Both CPUs have concurrent access to any drive on the DSSI.
 - A single system disk is accessed via two paths and can be served to all satellites by either CPU.
 - If either CPU fails, satellites will access their disks through the remaining CPU.

- In a CI-based or mixed-interconnect configuration that includes HSC50 or HSC70 disk servers, running the VAX Volume Shadowing product will:
 - Increase disk availability
 - Provide performance advantages of multiple system and data disks
 - Preserve the ease of management of a single system disk
- The optimal VAXcluster configuration for any computing environment is based on trade-offs of cost, performance and functionality. The performance of a VAXcluster configuration is influenced by the following factors:
 - CPU model of the members
 - Type of interconnect (CI and/or Ethernet) and adapters
 - Disk I/O capacity and access time
 - Number of disks being served
 - Applications in use
 - Ethernet utilization

HARDWARE REQUIREMENTS

Processors Supported:

VAX: VAX 6210, VAX 6220, VAX 6230, VAX 6240, VAX 8200, VAX 8250, VAX 8300, VAX 8350, VAX 8500, VAX 8530, VAX 8550, VAX 8600, VAX 8650, VAX 8700, VAX 8800, VAX 8810, VAX 8820, VAX 8830, VAX 8840

VAX-11/750, VAX-11/780, VAX-11/785

MicroVAX: MicroVAX II, MicroVAX 2000, MicroVAX 3300, MicroVAX 3400, MicroVAX 3500, MicroVAX 3600

VAXstation: VAXstation II, VAXstation 2000, VAXstation 3200, VAXstation 3500, VAXstation 8000

VAXserver: VAXserver 3300, VAXserver 3400, VAXserver 3500, VAXserver 3600, VAXserver 3602, VAXserver 6210, VAXserver 6220

Preconfigured VAXcluster Systems:

VAX: VAX 6200, VAX 8842, VAX 8974, VAX 8978

Processors Not Supported:

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

MicroVAX: MicroVAX I

VAXstation: VAXstation I

Processor Restrictions:

VAXcluster hardware requirements are dependent upon the environment that is configured.

CI-based VAXcluster systems require Computer Interconnect cables, adapter ports in each CPU, and a Star Coupler device.

VAXcluster Software supports the following Ethernet adapters:

- DESVA
- DEQNA *
- DELQA
- DEUNA
- DELUA
- DEBNA

* All systems utilizing a DEQNA must operate with software data checking enabled. Because AUTOGEN automatically sets the correct parameter, no system management intervention is required.

VAXcluster Software supports the following CI port adapters:

- CIBCI
- CIBCA-A / CIBCA-B
- CI750
- CI780

Boot servers must have a minimum of an RF30 acting as a system disk.

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

OPTIONAL HARDWARE

CI-based and mixed-interconnect clusters may use HSC50 or HSC70 intelligent mass storage controllers. These controllers offload disk management functions from host systems and provide several additional benefits, including:

- Increased I/O throughput
- Multiple host access to storage
- Volume shadowing support

A Computer Interconnect Star Coupler Expander (CISCE) may be added to a CI-based or mixed-interconnect configuration to increase the total number of CI-connected nodes to 24 (VAX nodes and HSCs). The maximum number of VAX nodes allowed in a CI-based configuration utilizing a CISCE is 16. The maximum number of VAX nodes allowed in a mixed-interconnect configuration utilizing a CISCE is 42 (maximum of 16 CI-connected VAX nodes).

The HSC-series Controller

- Each HSC50 supports a maximum of six channels for disks and/or tapes.
- Each HSC70 supports a maximum of eight channels for disks and/or tapes.
- Each disk channel supports a maximum of four disk drives and each tape channel supports a maximum of four tape subsystems (with four tape drives each).

The VMS Operating System supports a maximum of 12 tape drives per HSC-series controller.

- Version 5.1 of the VMS Operating System supports dual-pathing and failover of disks between pairs of HSCs, between local controllers, and between disk servers. 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. Failover of disk drives dual pathed between HSC controllers, between local adapters, or between disk servers helps to provide high availability. When a disk server fails in configurations that include multiple servers, satellite access to disks will fail over to another server. The failover process is transparent to applications.
- Dual-pathing of TA-series tape drives between two HSCs is also supported by the VMS operating system.

Note: HSC Software, at least at V3.5 is required to support dual-pathing of TA drives.

SOFTWARE REQUIREMENTS

- VMS Operating System V5.1

Note: Version 5.0, 5.0-n and 5.1 of the VMS Operating System may coexist within a VAXcluster system. However, DIGITAL recommends that the latest release of the operating system be installed on each VAX node in the cluster.

For CI-based VAXcluster configurations only, version 5.0, 5.0-n and 5.1 of the VMS Operating System offer a mode of operation that is compatible with Version 4.7.

- DECnet-VAX V5.1

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

DECnet-VAX software supports the ability to access some or all nodes in a VAXcluster system using a separate alias node address, while retaining the ability to address each node in the cluster individually. Not all network objects may be accessed using this mechanism. This feature requires that at least one node operate as a full function node.

DECnet-VAX software is also needed if process-to-process communication is desired between processes on different VAXcluster nodes under program control. The MONITOR utility requires DECnet-VAX for some MONITOR activities.

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

OPTIONAL SOFTWARE

Refer to the VAXcluster Support section of the Software Product Descriptions for optional VMS software products.

Optional products that are particularly useful in VAXcluster systems include VAX Volume Shadowing (SPD 27.29.xx), VAX Performance Advisor (SPD 27.71.xx) and VAXcluster Console System (SPD 27.46.xx).

ORDERING INFORMATION

Software License: QL-VBRA*-AA

Services: QT-VBRA*-**

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

SOFTWARE LICENSING

License Options

A license is required on each CPU in a VAXcluster system, regardless of configuration type.

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

Licensing Management Facility Support

This layered product supports the VMS License Management Facility.

License units for this product are allocated on a CPU-capacity basis.

For more information on the License Management Facility, refer to the VMS Operating System Software Product Description (SPD 25.01.xx) or the License Management Facility manual of 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.