Software Product Description

PRODUCT NAME: HP Galaxy Software Architecture on OpenVMS Alpha, Version 1.1
SPD 70.44.09

Note:

This SPD describes the HP Galaxy Software Architecture on OpenVMS, which is available as a separately licensed System Integrated Product (SIP).

Throughout this document, **OpenVMS Galaxy** and **Galaxy** refer to the **Galaxy Software Architecture on OpenVMS**. The term **instance** refers to a copy of the OpenVMS Alpha operating system.

DESCRIPTION

OpenVMS Galaxy enables multiple instances of OpenVMS to execute cooperatively in a single computer, giving customers the ability to manage unpredictable, variable, or growing workloads.

With OpenVMS Galaxy, many processors and other physical resources are partitioned in order to run multiple instances of operating systems. Each instance has assigned CPUs, memory, and I/O. The instances share a part of the memory, and CPUs can be reassigned from one instance to another while the system is running. This computing environment can be dynamically adapted to meet changing application needs and workload demands.

Software logically **partitions** systems by assigning CPUs, memory, and I/O ports to individual instances of the OpenVMS operating system. This partitioning, which a system manager directs, is a software function; no hardware boundaries are required. Each individual instance has the resources it needs to execute independently. An OpenVMS Galaxy environment is **adaptive** in that resources such as CPUs can be dynamically reassigned to different instances of OpenVMS.

Memory is logically partitioned into private and shared sections. Each instance of the operating system has its own private memory, which means that no other instance writes to those physical pages. Some of the shared memory is available for instances of OpenVMS to communicate with one another, and the rest of the shared memory is available for application data.

An OpenVMS Galaxy has a highly scalable I/O subsystem because the system contains multiple primary CPUs — one in each instance. In addition, OpenVMS currently has features for distributing some I/O to secondary CPUs in a symmetric multiprocessing (SMP) system.

CPUs within an OpenVMS Galaxy are allocated to instances.

In an OpenVMS Galaxy environment, the **console firmware** plays a critical role in partitioning hardware resources. It maintains the permanent configuration in NVRAM and the running configuration in memory. The console provides each instance of the OpenVMS operating system with a pointer to the running configuration data.

The console performs power-up self-tests, initializes hardware, initiates system booting, and performs I/O services during system booting and shutdown. The console program also provides run-time services to the operating system for console terminal I/O, retrieval of environment variables, NVRAM saving, and other services.

An OpenVMS Galaxy computing environment lets customers decide how much cooperation exists between instances in a single computer system.

In a **shared-nothing** computing model, the instances do not share any resources; operations are isolated from one another.

In a **shared-partial** computing model, the instances share some resources and cooperate in a limited way.

In a **shared-everything** model, the instances cooperate fully and share all available resources.

OpenVMS Galaxy leverages proven OpenVMS Cluster, SMP, and performance capabilities to offer greater levels of performance, scalability, availability, and flexibility.

By running multiple instances of OpenVMS in a single computer, an OpenVMS Galaxy computing environment gives you quantum improvements in:

- Compatibility—Existing applications run without changes.
- Availability—Presents opportunities to upgrade software and expand system capacity without downtime.
- Scalability—Offers scaling alternatives that improve performance of SMP and cluster environments.
- Adaptability—Physical resources can be dynamically reassigned to meet changing workload demands.
- Cost of ownership—Fewer computer systems reduce system management requirements, floor space, and more.

For companies looking to improve their ability to manage unpredictable, variable, or growing IT workloads, OpenVMS Galaxy technology provides a flexible way to dynamically reconfigure and manage system resources.

An OpenVMS Galaxy computing environment is ideal for high-availability applications, such as:

- Database servers
- Transaction processing systems
- Data warehousing
- Data mining
- Internet servers

Software and Hardware Components

An OpenVMS Galaxy computing environment is comprised of the following:

- OpenVMS Alpha operating system
- AlphaServer Console Firmware
- Supported hardware

For more information about these components, see the Software Requirements and Hardware Requirements sections of this SPD.

FEATURES

You can create an OpenVMS Galaxy computing environment that allows you to:

- Reassign CPUs between instances.
- Perform independent booting and shutdown of instances.
- Use shared memory for interinstance communication.
- Create a shared memory RAMdisk with HP DECram for OpenVMS Alpha Version 3.0 or later.
- Cluster instances within an OpenVMS Galaxy using the Shared Memory Cluster Interconnect (SMCI).
- Cluster instances with non-Galaxy systems.
- Create applications using OpenVMS Galaxy application programming interfaces (APIs) for resource management, event notification, locking for synchronization, and shared memory for global sections.
- Use the Galaxy Configuration Utility (GCU) to view and control the OpenVMS Galaxy environment.
- Run a single-instance OpenVMS Galaxy on any Alpha system for application development.

The following sections describe some of these features in more detail.

Galaxy Configuration Utility

The Galaxy Configuration Utility is a DECwindows Motif application that allows system managers to configure and manage an OpenVMS Galaxy system from a single workstation window.

Using the GCU, system managers can:

- Display the active Galaxy configuration.
- Reassign resources among Galaxy instances.
- View resource-specific characteristics.
- Shut down or reboot one or more Galaxy instances.
- Invoke additional management tools.
- Create and engage Galaxy configuration models.
- Create a single-instance Galaxy on any Alpha system (for software development on non-Galaxy hardware platforms).
- View the online Galaxy documentation.
- Determine hot-swap characteristics of the current hardware platform.

Shared Memory Cluster Interconnect

The Shared Memory Cluster Interconnect (SMCI) is a System Communications Services (SCS) port for communications between Galaxy instances. When an OpenVMS instance is booted as both a Galaxy and as an OpenVMS Cluster member, the SMCI driver is loaded. This SCS port driver communicates with other cluster instances in the same Galaxy through shared memory. This capability provides one of the major performance benefits of the OpenVMS Galaxy Software Architecture. The ability to communicate to another clustered instance through shared memory provides dramatic performance benefits over traditional cluster interconnects.

Local Area Network (LAN) Shared Memory Device Driver

Local Area Network (LAN) communications between OpenVMS Galaxy instances are supported by the Ethernet LAN shared memory driver. This LAN driver communicates to other instances in the same OpenVMS Galaxy system through shared memory. Communicating with other instances through shared memory provides performance benefits over traditional LANs.

RAMdisk in Shared Memory

In DECram for OpenVMS Alpha Version 3.0 or later, DECram's capability is extended to use OpenVMS Galaxy shared memory to create an OpenVMS shared memory disk. This allows users to take advantage of OpenVMS Galaxy shared memory without modifications to any of their applications.

Application Programming Interfaces

- Locks for synchronization
- Event notification
- Shared memory global sections
- Configuration information
- CPU management

For more information about OpenVMS Galaxy APIs, refer to the HP *OpenVMS Alpha Partitioning and Galaxy Guide*.

Single-Instance Galaxy Configuration

A single-instance Galaxy is for non-Galaxy platforms, that is, those without a Galaxy console. Galaxy configuration data, which is normally provided by console firmware, is instead created in a file. By setting the system parameter GALAXY to 1, SYSBOOT reads the file into memory and the system boots as a single-instance Galaxy, complete with shared memory, Galaxy system services, and even self-reassignment of CPUs. This can be done on any Alpha platform.

Single-instance Galaxy configurations will run on any Alpha workstations or servers (even laptops) running OpenVMS Version 7.2–1 or higher. This capability allows early adopters to evaluate OpenVMS Galaxy features and, most important, to develop and test Galaxy-aware applications without incurring the expense of setting up a full-scale Galaxy platform.

Because a single-instance Galaxy is not an emulator—it is real Galaxy code—applications developed on a single-instance Galaxy will run on multiple-instance configurations.

SOFTWARE REQUIREMENTS

OpenVMS Galaxy configurations require OpenVMS Alpha Version 7.2–1 or higher.

OpenVMS Galaxy console firmware for the supported AlphaServers is located on the Alpha Systems Firmware Update CD–ROM.

HARDWARE REQUIREMENTS

Configuring an OpenVMS Galaxy computing environment requires:

- One I/O module for each partition
- At least one processor module for each partition
- A dedicated serial console port for each partition
- Sufficient memory for operating system and required applications

Optional Hardware

Customers might want to use the following optional hardware:

- I/O expansion from an I/O module
- I/O adapters for network, storage, or traditional cluster interconnects
- Additional CPUs for SMP instances
- Additional memory for Galaxy-wide global sections

GCU Hardware Requirements

Hewlett-Packard recommends an Alpha or VAX workstation running DECwindows or a Windows NT workstation with an X terminal emulator as a display device for the OpenVMS Galaxy Configuration Utility.

SUPPORTED HARDWARE

AlphaServer ES47, ES80, and GS1280 Galaxy Configuration

OpenVMS Alpha Version 7.3–1 or higher supports the following maximum configurations on these AlphaServer systems:

- AlphaServer ES47 M4 (2 instances)
- AlphaServer ES80 M8 (4 instances)
- AlphaServer GS1280 M8 (2 instances)
- AlphaServer GS1280 M16 (4 instances)
- AlphaServer GS1280 M8 (4 instances)
- AlphaServer GS1280 M16 (8 instances)
- AlphaServer GS1280 M32 (8 instances)
- AlphaServer GS1280 M64 (16 instances*)

*Note: M64 requires at least 2 hard partitions with a maximum of 8 instances each.

Configuration Rules

Note: The AlphaServer ES80 and GS1280 systems can be configured with one or more hard partitions. OpenVMS Galaxy can be installed on each hard partition to configure one or more soft partitions.

Each soft partition must have at least:

- 1 Console interface per CPU
- Each CPU can be a primary (only 1 primary per 2P duo)
- 1 I/O port (IO7)
- Minimum required memory per CPU

AlphaServer GS80/160/320 Galaxy Configuration

OpenVMS Alpha Version 7.2–1H1 or higher supports the following maximum configurations on these AlphaServer systems:

- AlphaServer GS80 (2 instances)
- AlphaServer GS160 (4 instances)
- AlphaServer GS320 (8 instances)

Configuration Rules

Note: The AlphaServer GS160, and GS320 systems can be configured with one or more hard partitions. OpenVMS Galaxy can be installed on each hard partition to configure one or more soft partitions.

Each soft partition must have at least:

- A standard COM1 UART console line
- A master PCI drawer
- An I/O module
- One CPU module
- One memory module

AlphaServer ES40 Galaxy Configuration

AlphaServer ES40 systems can support a maximum of two instances of OpenVMS.

AlphaServer ES40 Clock

An AlphaServer ES40 has one clock. For an OpenVMS Galaxy, this means that you cannot run the two instances at different times. Also, the SET TIME command affects both instances. Note that this may not become evident until a number of hours have passed.

Console Ports

On a rack-mounted system:
COM1 (lower) is the console port for instance 0.
COM2 (upper) is the console port for instance 1.

On a pedestal system:
COM1 (left) is the console port for instance 0.
COM2 (right) is the console port for instance 1.

Unlike creating an OpenVMS Galaxy on an AlphaServer 8400 system, you do not need additional hardware for the second console. COM2 is used for this purpose.

CPUs

CPU0 must be the primary for instance 0.
CPU1 must be the primary for instance 1.
CPUs 2 and 3 are optional secondary CPUs that can be migrated.

I/O Adapters

On a rack-mounted system:
PCI Hose 0 (PCI0) belongs to instance 0 (upper 4 PCI slots)
PCI Hose 1 (PCI1) belongs to instance 1 (lower 6 PCI slots)

On a pedestal system:
PCI Hose 0 (PCI0) belongs to instance 0 (right-hand slots)
PCI Hose 1 (PCI1) belongs to instance 1 (left-hand slots)

Note that PCI0 contains an embedded ISA controller.

Storage Controllers

You will need one storage controller (such as a KZPSA) per instance. For each instance, this can go to a separate StorageWorks box or to the same box for running as a SCSI cluster.

Network Cards

If each instance needs network access, a network card (such as a DE600) is required for each instance.

One card each goes in PCI0 and PCI1.

Memory Granularity Restrictions

Private memory must start on a 64-MB boundary.

Shared memory must start on an 8-MB boundary.

Instance 0 must have a multiple of 64-MB.

AlphaServer 8400/GS140 Galaxy Configuration

AlphaServer 8400 systems can support two or three instances of OpenVMS.

9 slots for:

I/O modules
Memory modules
Processor modules (2 CPUs per module)

Console line for each partition:

Standard UART for first partition
KFE72-DA for each additional partition

Example Configuration 1:
2 partitions, 8 CPUs, 12 GB memory

9 slots allocated as follows:

Two I/O modules
Four processor modules (2 CPUs each)
Three memory modules (4 GB each)

Example Configuration 2:
3 partitions, 8 CPUs, 8 GB memory

9 slots allocated as follows:

Three I/O modules
Four processor modules (2 CPUs each)
Two memory modules (4 GB each)

AlphaServer 8200/GS60/GS60E Galaxy Configuration

AlphaServer 8200, GS60, and GS60E systems can support one possible OpenVMS Galaxy configuration (two instances only).

5 slots for:

Two processor modules (two CPUs each)
Two I/O modules
One memory module

The AlphaServer GS140, GS60, GS60E, 8400, and 8200 systems provide a built-in UART for the first console line. Each additional console requires a module set.

The only hardware required for Galaxy operation that is not in the typical AlphaServer GS140, GS60, GS60E, 8400, and 8200 configuration is the KFE72-DA console subsystem.

The KFE72-DA module set is the set of EISA bus modules that establish a console port.

XMI Bus Support

The XMI bus is supported only on the first instance (instance 0) of a Galaxy configuration in an AlphaServer 8400 system.

AlphaServer 8400 systems can support only one DWLM-AA XMI PIU subsystem cage for all XMI devices. Note that the DWLM-AA takes up considerable space in the system because an I/O bulkhead is required on the back of the system to connect all XMI devices to the system. This allows only two additional DWLPB PCI PIUs in the system.

EISA Bus Support

The EISA bus is supported only on the first instance (instance 0) of a Galaxy configuration. Due to the design of all EISA options, they must always be on instance 0 of the system. A KFE70 must be used in the first instance for any EISA devices in the Galaxy system.

All EISA devices must be on instance 0. No EISA devices are supported on any other instance in a Galaxy system.

A KFE72-DA installed in other instances provides console connection only and cannot be used for other EISA devices.

AlphaServer 4100 Configuration

To create an OpenVMS Galaxy on an AlphaServer 4100 system, you must observe the following configuration and hardware requirements:

Two-Instance Maximum

You can run a maximum of two instances of OpenVMS on an AlphaServer 4100.

Console Firmware

You must have AlphaServer 4100 console firmware, which is located on the Alpha Systems Firmware Update Version 5.4 CD–ROM that is included with the OpenVMS Alpha Version 7.2–1 CD–ROM package.

AlphaServer 4100 Clock

An AlphaServer 4100 has one clock. For an OpenVMS Galaxy, this means that you cannot run the two instances at different times. Also, the SET TIME command affects both instances. Note that this may not become evident until a number of hours have passed.

Console Ports

COM1 (upper) is the console port for instance 0.
COM2 (lower) is the console port for instance 1.

Unlike creating an OpenVMS Galaxy on an AlphaServer 8400, you do not need additional hardware for the second console. COM2 is used for this purpose.

CPUs

CPU0 must be the primary for instance 0.
CPU1 must be the primary for instance 1.
CPUs 2 and 3 are optional secondary CPUs that can be migrated.

I/O Adapters

The four lower PCI slots belong to IOD0, which is the I/O adapter for instance 0.
The four upper PCI slots belong to IOD1, which is the I/O adapter for instance 1.

Storage Controllers

You will need two storage controllers (such as KZPSAs). These can go to separate StorageWorks boxes or to the same box for running as a SCSI cluster. One controller each goes in IOD0 and IOD1.

Network Cards

If each instance needs network access, a network card (such as a DE500) is required for each instance.

One card each goes in IOD0 and IOD1.

Physical Memory

Because OpenVMS Galaxy on an AlphaServer 4100 does not support memory holes, physical memory for an OpenVMS Galaxy environment must be contiguous. To achieve this on an AlphaServer 4100, one of the following must be true:

- All memory modules must be the same size (for example, 1 GB).
- If two sizes are present, only one module can be a smaller size. You must put the larger modules into the lower-numbered slots.

Note: OpenVMS Alpha Version 7.3–2 is the last version to support OpenVMS Galaxy on the AlphaServer 4100 and ES40.

LICENSING REQUIREMENTS

The Galaxy Software Architecture on OpenVMS (OpenVMS Galaxy) is a System Integrated Product (SIP): the OpenVMS Galaxy code is integrated and delivered with the OpenVMS operating system.

The License Management Facility (LMF) Product Authorization Keys (PAKs) representing OpenVMS Galaxy licenses allow you to access and use OpenVMS Galaxy software. For more information about the location of the PAKs available with OpenVMS Alpha Version 7.3–1 or higher, see the corresponding *Guide to OpenVMS CD-ROMs*.

The following list summarizes OpenVMS Galaxy licensing requirements:

- One OpenVMS Operating System License for a Galaxy system
- One SMP Extension License for each CPU after the first CPU
- One OpenVMS Galaxy License for each CPU in a Galaxy system
- No changes to how HP layered products are licensed:

  One capacity license per system
  One user license per use

The following sections describe these requirements in more detail.

OpenVMS Operating System License

When an AlphaServer system is configured as an OpenVMS Galaxy system, there are no changes in how a system is licensed for the OpenVMS operating system.

One OpenVMS Base License is required for the Galaxy system, plus one SMP Extension License for each CPU after the first CPU.

OpenVMS Galaxy License

In order to create and run multiple instances, one OpenVMS Galaxy License is required for each CPU in a Galaxy system.

License rights for running a single-instance Galaxy on any Alpha system are provided by the OpenVMS Base License.

OpenVMS Layered Products License

HP software products that run on OpenVMS Galaxy configurations continue to use standard license types: Traditional, Concurrent Use, and Personal Use.

- One Traditional capacity License will continue to license the system, regardless of the number of instances. The license is based on the system class of the hardware system.
- Concurrent Use Licenses will continue to license one concurrent use of the product.
- Personal Use Licenses will continue to license one named user on the system.

Clustering OpenVMS Galaxy Instances

Instances in an OpenVMS Galaxy computing environment can be clustered with other instances in a single system, with instances in other Galaxy systems, or with non-Galaxy systems. Each type of clustering has different licensing requirements, as described in the following sections.

Clustering within a Galaxy System

In an OpenVMS Galaxy computing environment, instances can be clustered with other instances within a Galaxy system. Clustered instances use the shared-memory cluster interconnect to communicate with each other.

The licensing and functionality for clustering within a Galaxy system is provided under the OpenVMS Galaxy License.

Clustering Outside a Galaxy System

Instances in an OpenVMS Galaxy computing environment can be clustered with instances in another OpenVMS Galaxy system or with cluster nodes in non-Galaxy systems. Instances clustered outside of a Galaxy system use traditional cluster interconnects.

Each system that is clustered with another system must be licensed for OpenVMS Cluster Software. Clustering outside the OpenVMS Galaxy system is not covered by the OpenVMS Galaxy License.

DISTRIBUTION MEDIA

The Galaxy Software Architecture on OpenVMS Alpha is supplied on the same distribution media as the OpenVMS operating system. For more information, refer to the *HP OpenVMS Operating System for Alpha Version 7.3-1 and 7.3-2 and VAX Version 7.3 Software Product Description* (SPD 25.01.xx) or the *HP OpenVMS for Integrity Servers and HP OpenVMS Alpha Version 8.3 Operating Systems Software Product Description* (SPD 82.35.xx),

ORDERING INFORMATION

Software Licenses

QL–66XAA–3B	Galaxy 1 CPU License
QL–66XAA–3C	Galaxy 2 CPU License
QL–66XAA–3D	Galaxy 4 CPU License
QL–66XAA–3E	Galaxy 8 CPU License
QL–66XAA–3F	Galaxy 16 CPU License

Documentation

The *OpenVMS Alpha Partitioning and Galaxy Guide* describes how to create, manage, and use an OpenVMS Galaxy computing environment. This book includes:

- OpenVMS Galaxy hardware and configuration requirements
- Procedures for creating OpenVMS Galaxy computing environments on OpenVMS AlphaServers listed in the supported hardware section above.
- Complete details about how to use all of the OpenVMS Galaxy features and capabilities available in OpenVMS Alpha Versions 7.2–1 through 8.3.

The *OpenVMS Alpha Partitioning and Galaxy Guide* is available on the OpenVMS documentation CD-ROM in HTML and PDF formats.

SOFTWARE LICENSING

This software is furnished under the licensing provisions of Hewlett-Packard's Standard Terms and Conditions.

License units for OpenVMS Galaxy are allocated on a CPU-capacity basis. An OpenVMS Galaxy license is required for each CPU within a system.

License Management Facility Support

This system integrated product supports the OpenVMS License Management Facility.

For more information on the License Management Facility, refer to the *HP OpenVMS Operating System for Alpha Version 7.3-1 and 7.3-2 and VAX Version 7.3 Software Product Description* (SPD 25.01.xx) or the *HP OpenVMS for Integrity Servers and HP OpenVMS Alpha Version 8.2 Operating Systems Software Product Description* (SPD 82.35.xx), or the appropriate operating system documentation set.

For more information about OpenVMS licensing terms and policies, contact your local HP sales office, or find HP software licensing information on the World Wide Web at:

http://h18000.www1.hp.com/products/software/info/terms/swl_sld.html

SOFTWARE PRODUCT SERVICES

A variety of service options are available from HP. For more information, contact your local HP account representative or distributor. Information is also available from http://www.hp.com/hps/software.

SOFTWARE WARRANTY

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

© 2006 Hewlett-Packard Development Company, L.P.

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

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