Breakthrough OpenVMS Galaxy Software

Looking to improve your ability to manage unpredictable, variable, or growing IT workloads? Then meet the high-end server future! OpenVMS Galaxy technology provides a flexible way to dynamically reconfigure and manage your system resources. An integrated hardware and software solution, OpenVMS Galaxy lets you perform tasks such as reassigning individual CPUs through a simple drag-and-drop procedure or automatically under program control.

Speed

Total cost of ownership

Availability

Reliability

Scalability

Embracing large multi-CPU systems via Adaptive Partitioned Multiprocessing (APMP), OpenVMS Galaxy leverages your existing cluster technology, overcomes software SMP and MPP limitations, and offers a viable commercial machine for unlimited high-end computing.

An OpenVMS Galaxy computing environment is ideal for applications such as:

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

Six great ways OpenVMS Galaxy adds value to IT investments

OpenVMS Galaxy is innovation in action — another example of Compaq's commitment to customers' enterprise computing needs. With it you can:

- **Dynamically reassign system resources, mapping compute power to applications on an as-needed basis**
- **Get more out of a single system (AlphaServer 8000 and 4100 systems) by running multiple coordinated instances of OpenVMS inside**
- **Expand cluster capability — inside and outside the system**
- **Gain scaling alternatives that improve performance of SMP and cluster environments**
- **Use memory as a cluster interconnect — and measure transfers in nanoseconds**
- **Upgrade software and perform maintenance, without downtime**

... “We are committed to maintaining our lead in continuously-available, business-critical solutions by continuing our investment in OpenVMS…”

— Eckhard Pfeiffer, President and CEO
Compaq Computer Corporation

Global 24x365 support for enterprise systems

Compaq is one of the few organizations capable of delivering consistent quality services around the world. We have local service offices in more than 100 countries for the support of global environments — around the clock and around the world.

Take the next step

To put OpenVMS Galaxy technology to work balancing resources in your enterprise, get connected with your OpenVMS ambassador at 1-800-344-4825 or call your local Compaq office. Visit our web page at: **www.compaq.com/openvms**

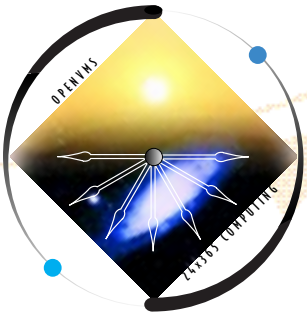

OpenVMS Galaxy features At-A-Glance

With OpenVMS Alpha V7.2, you can create an OpenVMS environment that allows you to:

- Run three instances of OpenVMS on an AlphaServer 8400
- Run two instances of OpenVMS on an AlphaServer 4100
- Run two instances of OpenVMS on an AlphaServer 8200
- Reassign CPUs between instances
- Perform independent booting and shutdown
- Use shared memory for inter-instance communication
- Cluster instances with non-galaxy systems
- Cluster instances within an OpenVMS Galaxy using the shared memory cluster interconnect
- Create applications using OpenVMS Galaxy APIs for resource management, event notification, locking for synchronization, and shared memory for global sections
- Use the Galaxy Configuration Utility to view and control the OpenVMS Galaxy environment
- Run a single-instance OpenVMS Galaxy on any Alpha system for application development

www.compaq.com/openvms

Printed in the U.S.A. EC-F9574-93 REL. #00/00 00 00 00.0 Copyright 1998 Compaq Computer Corporation. All Rights Reserved.

COMPAQ

Compaq Galaxy Software Architecture

Introducing the next generation of OpenVMS on Alpha

In a perfect world, you'd never have to take your system down for *any* reason. You'd have the ability to balance your workload on-the-fly — accommodating multiple applications and changing resource demands. And you'd be able to do the same amount of work with fewer computers. A tall order? A simple answer: Compaq Galaxy Software Architecture on OpenVMS™ Alpha. OpenVMS Galaxy leverages proven OpenVMS technologies — bringing you quantum improvements in flexibility and cost-of-ownership.

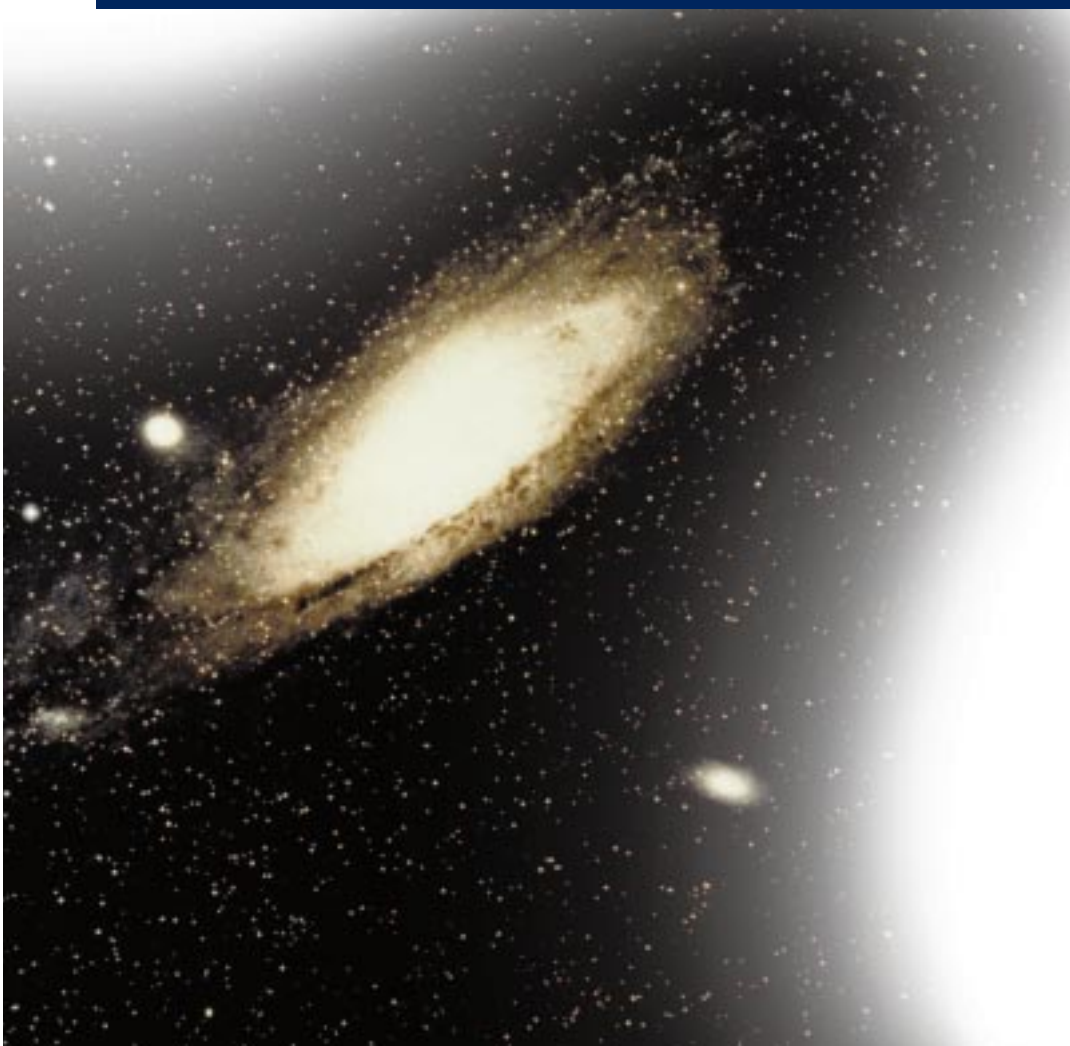

Benefits

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 shared and dynamically reassigned
- **Cost of ownership —** Fewer computer systems reduce system management requirements, floor space, and more

What is an OpenVMS Galaxy?

The Basic Concept

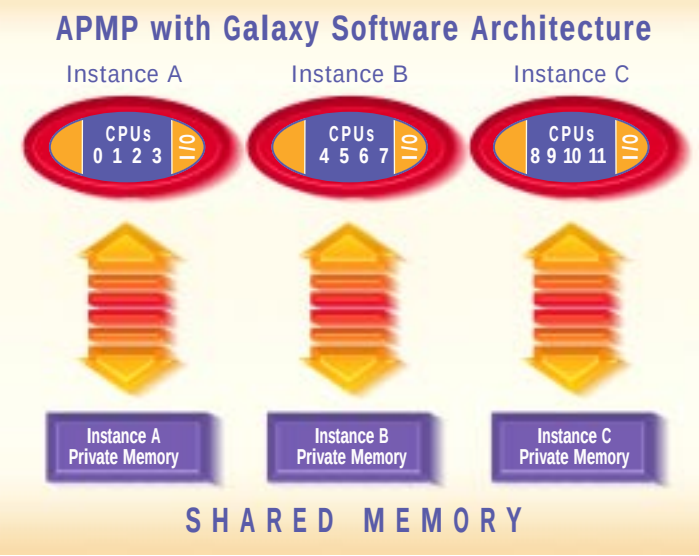

Adaptive Partitioned Multi-Processing (APMP) is a new model of computing in which multiple instances of operating systems execute cooperatively in a single computer. The Compaq Galaxy Software Architecture on OpenVMS is the complementary software that leverages APMP.

You asked for it!

Compaq is developing an evolution of OpenVMS functionality that includes APMP, a new model of computing, which allows multiple instances of OpenVMS to run cooperatively within a single computer. This computing environment further extends the unlimited high end of OpenVMS — providing a premier server for enterprise computing that will take maximum advantage of high-end systems of the future.

Early in our development efforts, we asked many of our OpenVMS customers and business partners, "What do you want and need in computing?", "What do you need in future enhancements of OpenVMS?" Your answers focused on three key themes: expansion, business flexibility, and consolidation.

You also said that you didn't want to wait for future generations of AlphaServers — you wanted the Compaq Galaxy Architecture on OpenVMS *now*. You gave us many examples of where aspects of the Compaq Galaxy design could provide improvements on existing platforms that would make a significant difference to your business — especially in the area of cost-of-ownership.

With multiple instances of OpenVMS in memory at the same time, an OpenVMS Galaxy provides the following benefits to your organization:

Higher Availability

With OpenVMS Galaxy technology, you can upgrade both applications and operating system instances without disrupting your entire operation. With OpenVMS V7.2, you can upgrade your OpenVMS Galaxy instances one at a time, similar to OpenVMS Cluster rolling upgrades. This means you never have to take the entire system down.

The same is true if you need to AUTOGEN to re-establish parameter settings. If one instance is shut down or fails, all applications in the other instances in the OpenVMS Galaxy continue to run without interruption. Bottom line — your business operations never miss a beat!

Superior Scalability

An OpenVMS Galaxy has the potential for nearly linear scaling of any resource. The factor that usually limits scalability is the operating system overhead necessary to synchronize many resources in an SMP configuration. In other words, while your platform allows more physical resources, the software prevents you from making use of them.

OpenVMS Galaxy overcomes this limitation by offering multiple instances of the operating system. With multiple instances, any one instance simply has to synchronize access to a subset of the resources in the system — giving you the potential for linear scaling within each instance. This means that with judicious resource assignment, you can make maximum use of all your hardware resources. And that means significant cost-of-ownership benefits!

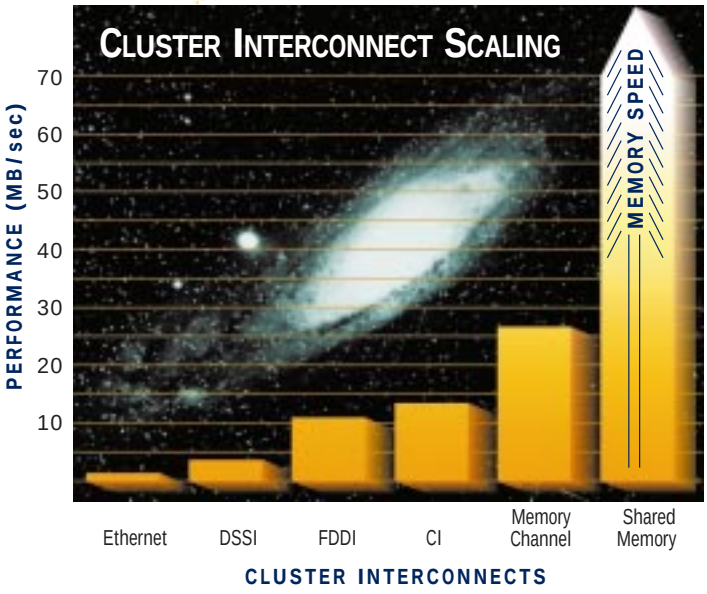

APMP and OpenVMS Galaxy, defined

Adaptive Partitioned Multi-Processing (APMP)

APMP is a new model of computing in which multiple instances of operating systems execute cooperatively in a single computer. With APMP, many processors (and other physical resources) are partitioned into multiple instances of operating systems. Each instance has assigned CPUs, memory, and I/O. The instances share a part of memory, and CPUs can be reassigned from one instance to another while the system is running. Many of the features, and hence the benefits of OpenVMS Galaxy technology, result directly from running multiple instances of the OpenVMS operating system in a single computer. This computing environment can be dynamically adapted to meet changing application needs and workload demands.

OpenVMS Galaxy Software Architecture

Compaq's first implementation of the APMP model of computing is the Compaq Galaxy Software Architecture on OpenVMS Alpha. The OpenVMS Galaxy Architecture provides the software structure and operating components that exploit the APMP computing model. An evolution in OpenVMS functionality, OpenVMS Galaxy leverages proven OpenVMS cluster, symmetric multiprocessing, and performance capabilities to offer greater levels of performance, scalability, availability and flexibility.

Single-instance OpenVMS Galaxy

This capability lets early adopters evaluate OpenVMS Galaxy features, and develop and test OpenVMS Galaxy-aware applications — without incurring the expense of setting up a full-scale OpenVMS Galaxy platform. This is for non-OpenVMS Galaxy platforms, that is, those without an OpenVMS Galaxy console. OpenVMS Galaxy configuration data, which is normally provided by console firmware, is instead created in a file on any Alpha platform. Because the single-instance OpenVMS Galaxy is not an emulator — it is real OpenVMS Galaxy code — applications you develop will run on multi-instance configurations.

Independent Instances

An OpenVMS instance that exists without sharing any resources is called an independent instance. One or more OpenVMS instances can execute without joining the sharing community. Such an instance does not access shared memory. An OpenVMS Galaxy may consist solely of independent instances — thus resembling traditional mainframe-style partitioning.

OpenVMS Galaxy also lets you use a trial-and-error method to evaluate resources. For example, you can reassign CPUs among instances of OpenVMS until you find the most effective combination of resources. And since all instances of OpenVMS and their applications continue to run as CPUs are reassigned, neither your users nor applications are interrupted.

High Adaptability

An OpenVMS Galaxy is highly adaptive because you can dynamically reassign computing resources to other instances of the operating system. Simply reassign processing power to the instances whose applications most need it. As that need changes, so can your configuration.

For example, if you know that resource demands change at certain times, you can write a command procedure to reassign CPUs to other instances of OpenVMS and submit it to a batch queue. You could do the same to manage system load characteristics.

Improved Cost-of-Ownership

An OpenVMS Galaxy presents opportunities to upgrade existing computers and expand their capacity, whether they are cluster members or independent systems, with a single computer running multiple instances of the operating system. Using fewer computers greatly reduces system management requirements, as well as floor space.

Compatibility

Your existing applications will run without changes on instances in an OpenVMS Galaxy. And your existing OpenVMS cluster applications will also run without changes on clustered instances in an OpenVMS Galaxy.

Galaxy Capabilities	Performance	Flexibility	Availability
Linear performance over a large # of CPUs (6+)	✶		
Shared memory for application usage across instances	✶	✶	✶
Multiple computing models (APMP)		✶	✶
Independent instances of operating systems		✶	✶
Dynamic reconfiguration on the fly (no reboot)		✶	✶
Upgrades		✶	✶
Scalable I/O	✶		
CPU Migration (Drag & Drop)		✶	✶
Can do "shared nothing" to "shared everything"	✶	✶	✶
Can implement clusters and failovers	✶	✶	✶

OpenVMS Galaxy Components

The Compaq Galaxy Software Architecture on OpenVMS includes the following hardware and software components:

AlphaServer Firmware Console — The console software 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, environment variable retrieval, NVRAM saving, and other services.

On an OpenVMS Galaxy, the console 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.

Shared Memory — Memory is logically partitioned into private and shared sections. Each operating system instance 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.

CPUs — CPUs within an OpenVMS Galaxy are allocated to instances.

I/O — An OpenVMS Galaxy has a highly scalable I/O subsystem because there are multiple, primary CPUs in the system — one in each instance. In addition, OpenVMS currently has features for distributing some I/O to secondary CPUs in an SMP system, and future advancements are in the works.

Choose the computing model that meets your needs

An OpenVMS Galaxy computing environment lets you decide how much cooperation exists between instances in a single computer system. In a "shared nothing" computing model, the instances do not share memory and they are not clustered with one another. In a "shared partial" computing model, the instances share memory but are not clustered with one another. In a "shared everything" model, the instances share memory and are clustered with one another. Which model you choose depends on your business needs.

Extend the power of your clusters

OpenVMS Cluster technology is the benchmark by which all other clusters are measured. Compaq has drawn on fifteen years of proven OpenVMS Cluster technology to facilitate communication among clustered instances within an OpenVMS Galaxy.

An OpenVMS Cluster is a set of OpenVMS operating systems, one per computer, that combine the processing power and storage capacity of multiple computers into a single, shared-everything environment. In contrast, an OpenVMS Galaxy is a set of OpenVMS operating systems in a single computer that communicates via shared memory.

OpenVMS Galaxy technology does not replace the OpenVMS cluster technology — rather, it's an enhancement to it. An OpenVMS Galaxy is a complete system. An instance of the operating system in an OpenVMS Galaxy can be clustered with other instances within the OpenVMS Galaxy system or with non-Galaxy systems. But although it can be added to an existing cluster, the OpenVMS Galaxy architecture is focused on a single system. The benefit is that a cluster-aware application running on instances of an OpenVMS Galaxy can take advantage of shared memory performance.