



*The OpenVMS Galaxy
Upgrade Advantage*

Consolidate

Discover the fast,
easy way to add
value to your
IT investments



Lower Costs



Flexibility

COMPAQ

Better answers

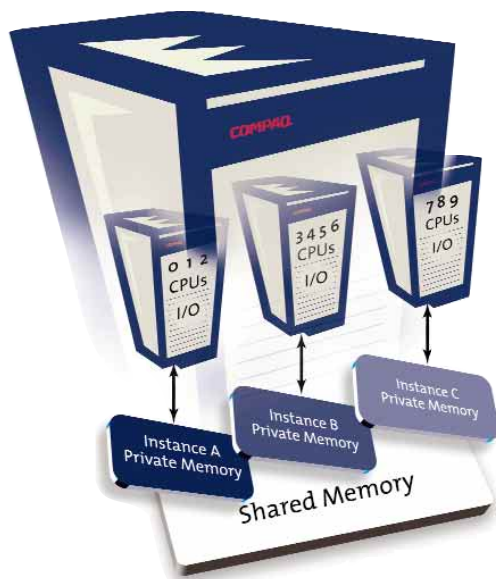
Investment protection

Leveraging proven *OpenVMS* technology

Compaq™ Galaxy software architecture on the *OpenVMS™ Alpha* OS lets enterprise IT infrastructures respond to business demands with unheard-of flexibility — at a fraction of the cost. The *OpenVMS Galaxy* architecture enhances the benefits of high availability and flexibility for which the *OpenVMS* OS has long been acclaimed.

It takes the best of cluster technology and multi-CPU technology — combining the two approaches into a single-system solution based on software partitioning. (See illustration.)

New *Galaxy* software architecture



With *OpenVMS Galaxy* architecture, a single system can run multiple copies, or instances, of the *OpenVMS* OS

Each instance has its own dedicated memory, which is part of the system's total memory component. Simply dedicate sections of memory to each instance, and then use the remaining memory as shared memory that all instances can access. This shared memory acts as a cluster interconnect within the system — thus achieving orders of magnitude improvement in performance. You can easily assign and reassign CPUs as needed, based on your application requirements.

Benefits of upgrading to the *OpenVMS Galaxy* architecture

- Allows you to do more with less by consolidating or redeploying servers
- Lowers cost of ownership
- Increases data and application availability
- Dynamically adapts system resources to meet changing business needs
- Lets you upgrade software and perform maintenance without downtime
- Runs existing applications without change
- Boosts performance by expanding cluster capabilities — inside and outside the system

Business flexibility



Add value to your IT investments

To keep ahead of the pack, you need to aggressively seek out new ways to make your IT infrastructure boost your business, making it more flexible and profitable. The *OpenVMS Galaxy* architecture offers many ways to give you an edge by adding value to your IT investments.

Dynamically reassign system resources — mapping compute power to applications on an as-needed basis

In an *OpenVMS Galaxy* environment, software is in total control of assigning and dynamically reassigning hardware resources. With a few clicks of a mouse, you can reassign computing resources to other instances of the operation system on the fly — while all applications remain running. If you know that resource demands change at certain times, you can write a command procedure to reassign CPUs to other instances of the *OpenVMS* OS and submit it to a batch queue. You could do the same to manage system load characteristics.

A stock exchange deployed the *OpenVMS Galaxy* architecture for this purpose. The exchange's IT infrastructure consisted of dedicated servers for batch operations, mission-critical order matching/trade processing, and data mining operations. The demands on their stock trading system were increasing, but it was already running at full capacity. Because of its ability to dynamically reassign CPUs between instances, the *OpenVMS Galaxy* architecture is providing the needed flexibility to handle increasing trade volumes — without interrupting trades. And this increased trading volume is boosting the stock exchange's revenues.

Get more out of single Compaq AlphaServer 8000 or 4100 systems by running multiple coordinated copies of the OpenVMS OS

Currently, the *OpenVMS Galaxy* architecture can be run on *AlphaServer*™ 8400 (3 instances) or *AlphaServer* 8200 or 4100 systems (2 instances). As a result, performance improves dramatically. For application development, a single-instance *OpenVMS Galaxy* architecture can be run on any *Alpha* system,

including laptops and workstations. This architecture will be supported on all new *Alpha* system releases — with the capability of running up to 24 or more instances on some servers.

The fact that you can distribute interrupts across instances provides many I/O configuration possibilities. You can partition a system's I/O workload so that certain I/O traffic is done on specific instances. As a result, performance improves dramatically.

An example of how the *OpenVMS Galaxy* architecture can boost performance comes from a medical research institution. The institution's strategic application receives bursts of input from data feeds that record brain waves, then performs number-crunching analysis on that data. With these two distinctly different tasks, it was difficult to tune their system for optimal performance. By partitioning with the *OpenVMS Galaxy* architecture, they were able to tune one instance to do all of the brain wave data collection, and tune the other instance to perform efficient analysis of the incoming data simultaneously. They achieved a very cost-effective solution on an *AlphaServer* 4100 system, achieving excellent performance in both dimensions and speeding up their operation by tuning the environment for each instance. This improved performance allows the institution to conduct research faster.

Gain scaling alternatives that improve the performance of SMP and cluster environments

Because the *OpenVMS Galaxy* architecture offers nearly linear scaling, you can add resources to match application requirements as your business needs grow or change. When you add a CPU to an *OpenVMS Galaxy* configuration, you can assign it to any instance of the *OpenVMS* OS. Typical SMP scaling issues do not restrict this architecture. You can define the number of *OpenVMS* OS instances, assign the number of CPUs in each instance, and control how they are used. You can also use a trial-and-error method to evaluate resources. For instance, you can move CPUs among instances until you find the most effective combination of resources. All instances of the *OpenVMS* OS and their applications continue to run as CPUs are reassigned, so your business never misses a beat.

Lower cost of ownership

For example, as your day-to-day business activity changes, each instance in the *OpenVMS Galaxy* architecture might experience changing peak workloads — and they may occur at varying times in the business day. This architecture can optimize the assignment of CPU resources so that the configuration changes as the demands for computing resources change. All of this can occur, of course, while your applications are still running to ensure that the optimal environment exists to meet the needs of your business.

Use shared memory as a cluster interconnect, and measure transfers in nanoseconds vs. 55 MB/sec cluster limits

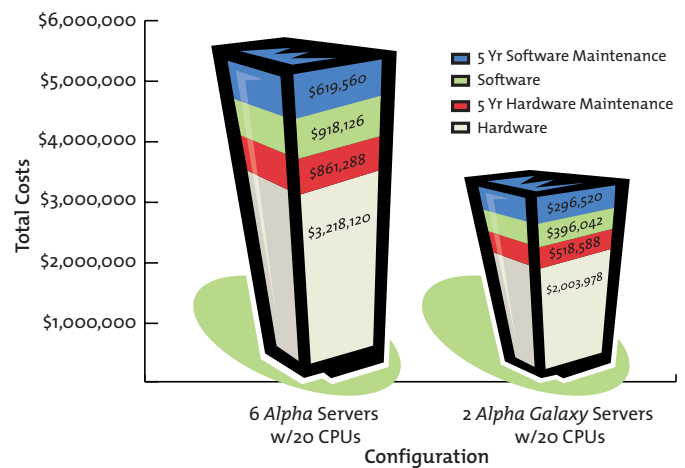
The *OpenVMS Galaxy* architecture leverages proven cluster, symmetric multiprocessing, and performance capabilities of the *OpenVMS OS* to offer greater levels of performance, scalability, and highly-available computing. One example of this leveraging is the shared memory cluster interconnect, which provides dramatic performance improvements over traditional cluster interconnects. Communications between instances of the *OpenVMS Galaxy* architecture take place through shared memory via the Shared Memory Cluster Interconnect (SMCI).

For example, a financial services bureau collects Windows NT® data feeds from worldwide sources and uses the *OpenVMS OS* for data manipulation and transmission. Customized data is delivered to clients' desktops, and I/O throughput is critical for providing the shortest time to delivering data. The *OpenVMS Galaxy* architecture provides the quickest time to data — which results in higher customer satisfaction and an increased number of customers.

Achieve improved total cost of ownership advantages

Another benefit of deploying this architecture is to achieve low cost of ownership. An *OpenVMS Galaxy* architecture gives you the opportunity to upgrade or re-deploy your existing computers and expand their capacity. Or you can consolidate some number of computers — whether they are cluster members or independent systems — with a single computer running multiple instances of the operating system. Having fewer computers significantly reduces maintenance costs, system management requirements, and floor space demands.

OpenVMS Galaxy architecture lets customers reduce system costs of ownership by 43% — plus streamline management and support requirements



The OpenVMS OS is ideally suited to workload and system consolidation, and is a cost-effective way to meet your growing performance requirements in the face of shrinking IT budgets. The OpenVMS Galaxy architecture can help you realize significant savings on operating system licenses and maintenance costs. Plus, it reduces system management headaches, and decreases floor space and power requirements.

*A VAX to Alpha migration will yield even greater results.

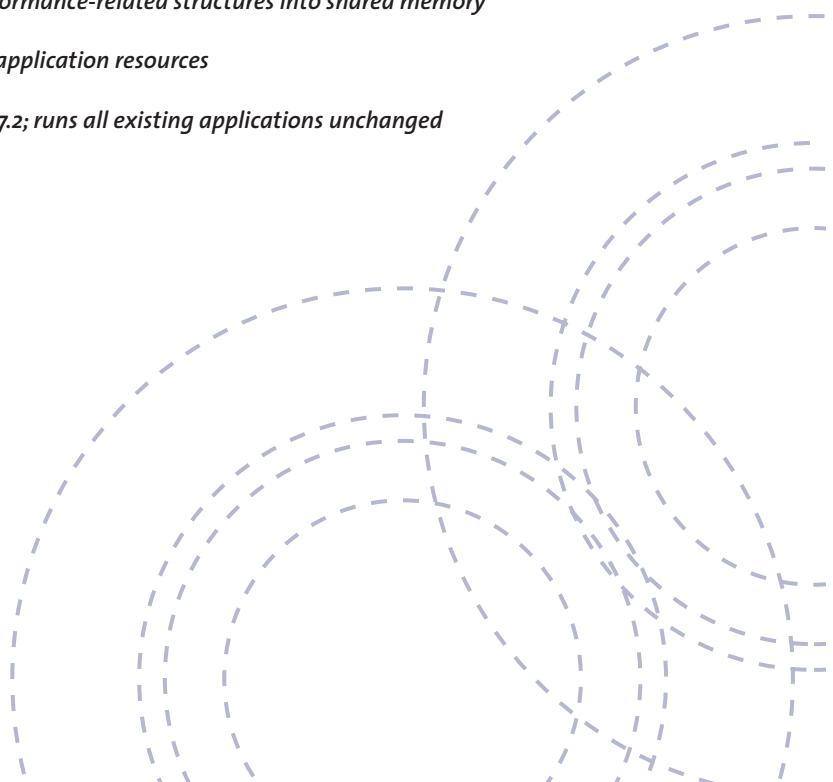


Quantum improvements in OpenVMS

When a technology comes along that lets you significantly expand the work you get out of your current systems, and consolidate or re-deploy your older systems to save on resources — while giving you the flexibility to instantly map your computer resources to your changing business needs — you make a move.

That technology is the Compaq *OpenVMS Galaxy* architecture, the next generation of the *OpenVMS* OS.

Features at a glance

- *Solves multi-CPU scaling limitations and cluster interconnect bottlenecks*
 - *Uses software to partition resources (CPUs, memory, and I/O) into separate OpenVMS OS instances*
 - *Employs shared memory as a cluster interconnect between instances of the operating system*
 - *Lets you move critical performance-related structures into shared memory*
 - *Provides management of application resources*
 - *Ships with OpenVMS OS V7.2; runs all existing applications unchanged*
- 

There's no better time to make a move

OpenVMS Galaxy ROI Opportunities at a Glance

Strategy	Return on Investment
Move applications off dedicated servers to more affordable systems	Put resources where you need them; get more value from each one
Upgrade existing computer(s)	Expand the capacity and extend the life of existing resources
Consolidate some number of computers, whether they are cluster members or independent systems, with a single computer running multiple instances of the operating system	Reduce system management headaches, and lower maintenance costs and requirements for floor space and power.
Cluster instances within an OpenVMS Galaxy system and with members of external clusters, or both	Use existing resources more flexibly; enjoy virtually unlimited scalability
Add an OpenVMS Galaxy environment to your existing resources, leveraging the ability to run different versions of OpenVMS OS simultaneously	Protect investments; build confidently with an eye on future technology enhancements

“The initial release provides a compelling proof that OpenVMS will live on as a viable, visible, and unlimited commercial solution for the installed base and new high-end customers alike. Endowed with the features and benefits provided by the Galaxy software architecture, the OpenVMS operating system will remain the premier server platform for bet-your-business, mission-critical, enterprise applications.”

Terry Shannon
Shannon Knows DEC, October 1998

Award-winning technology

In December 1998, the Compaq OpenVMS Galaxy architecture was awarded “Scene Setter” for the Serving Up Servers division of a computer industry competition for innovative technology. The event was sponsored by GIGA Information Group, one of the top IT analyst firms that provides proprietary research, analysis, and decision-support services to corporations investing in computer technology.

Take the next step

To learn more about upgrading to the OpenVMS Galaxy architecture, call 1-800-344-4825 in the U.S. and Canada. In other areas, contact your local sales representative or authorized business partner. For more detailed information, please visit our Web site at www.compaq.com/openvms

Compaq and the Compaq logo are copyrighted and are trademarks of Compaq Computer Corporation. Better answers is a service mark of Compaq Computer Corporation. AlphaServer and OpenVMS are registered with the U.S. Patent and Trademark office. Other products and company names mentioned herein may be trademarks and/or service marks of their respective owners.

Compaq believes the information in this publication is accurate as of its publication date; such information is subject to change without notice. Compaq is not responsible for any inadvertent errors.

Compaq conducts its business in a manner that conserves the environment and protects the safety and health of its employees, customers, and the community.

Printed in the U.S.A.
© Copyright 1999 Compaq Computer Corporation.
All rights reserved.
EC-F9883-71 Rel. #302/99 05 64 5.0

www.compaq.com/openvms

