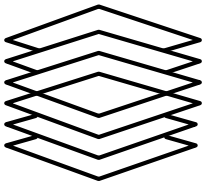

Enterprise Server Consolidation in the OpenVMS Environment

Harvey Hindin, Vice President,
High Availability Software and Hardware/Clusters Service

EXECUTIVE SUMMARY

Server consolidation offers massive benefits especially to companies with older servers and rapidly increasing workload. But server consolidation works best by focusing on the installed-base. Compaq's OpenVMS – the focus of this report – has many customers looking into server consolidation. A typical situation in this market segment is OpenVMS users with a number of VAX or older AlphaServer systems (of varying sizes) running multiple applications. OpenVMS on AlphaServer now offers greater clustering, availability, partitioning, and workload management than were available when the systems were installed. And these features are available at a reduced cost of ownership.

This report – one in a series on server consolidation from D. H. Brown Associates, Inc. (DHBA) – details a variety of consolidation methods. For example, all applications can be moved to a more powerful machine (faster, more memory, more storage) running on a single OpenVMS operating system instance on that large machine or a cluster. For the high-end AlphaServer GS-Series, it is possible to use Compaq's OpenVMS Galaxy software to set up multiple OpenVMS instances in partitions in either a single AlphaServer or AlphaServer cluster. The cluster architecture provides availability, scalability, and other advantages and, along with the Galaxy software provides major OpenVMS differentiators compared to other vendor offerings.

TYPICAL SERVER CONSOLIDATION BUSINESS DRIVERS

- Total Cost of Ownership reduction.
- Improved end user service level/maintenance of Service Level Agreements (SLAs).
- Floor space reduction.
- Consolidation of expensive datacenter real estate and redundant operations.

Physical (collocation) consolidation refers to the collocation of servers in a shared location. Typical benefits include economies of operational costs, such as system management, asset management, and operations. A second type, server data integration (storage) consolidation, involves standardization on a storage

platform. This includes large shared RAID systems and backup solutions, as well as the integration of servers and Storage Area Networks (SANs). Benefits span more efficient usage and management of disk resources, and better security. Type three is workload (application) consolidation and refers to the collocation of applications onto fewer, more powerful servers. Benefits include reduced hardware and software expenses, as well as additional reductions in operational costs.

COMPAQ'S OPENVMS ALTERNATIVES

Compaq's AlphaServer DS, ES, and GS Series systems running OpenVMS prove to be an attractive server consolidation target for existing OpenVMS environments This series of platforms provides high levels of performance, availability, and flexibility from the low-end through midrange and high-end server environments. The modular system architecture enables users to start with a system that meets immediate needs, but easily scales to larger systems as requirements change.

OpenVMS offers many benefits. For example, it allows alternate pathing with automatic failover (transparent to applications) for storage, network peripherals, and cluster interconnects with a single server. In a cluster environment, if a node fails unexpectedly, OpenVMS provides direct access to system resources (CPU, memory, disk) through other cluster nodes in the cluster. This access means that the only reconnect required would be the session connection itself, i.e., simply re-establish the connection on another node in the cluster and the process continues. There is no failover of disk resources, application restart, or failover scripts required to be written, because an OpenVMS Cluster can run applications on all systems at the same time. With OpenVMS Clusters it is also possible to install dynamic system patches while applications are running and available to users in a multi-node cluster.

In addition to clusters, on designated high-end AlphaServer systems, OpenVMS supports the online removal and replacement of failing CPU components.

Since the September 11 tragedy in the U.S., security and disaster recovery have risen to the top of the list of IT issues. One of the reasons why OpenVMS continues to be the choice for many mission critical environments is that it has long provided industry-leading disaster recovery capabilities with its unique active-active, load-balanced multi-site cluster technologies.

Three case studies of major organizations are presented in this report. One of the three is Verizon Information Services. To get the horsepower it needed, and to decrease TCO, Verizon replaced its production group of five VAX 7840 systems, each with 3.5 GB of memory, with three AlphaServer GS140 systems running OpenVMS 7.2, each with four processors and 14 GB of memory. The VAX to Alpha upgrade and concurrent consolidation resulted in cost savings reaching 40% in software costs, 25% in maintenance costs, and 20% less floor space. Moreover, it now takes six people to maintain the systems rather than eight.

There were still more benefits from the effort. For example, Verizon Information Services now enjoys a six fold improvement in response time (from 12 seconds to less than two seconds), and SLAs are now consistently met. Another bonus: productivity. Representatives now handle 20% more customers in the same time frame.

Compaq OpenVMS offers a highly integrated combination of software, VAX and/or AlphaServer computers, storage devices, and other shared peripherals that are suitable for server consolidation.

TABLE OF CONTENTS

INTRODUCTION AND OVERVIEW

Server consolidation in the OpenVMS environment is, for the most part, a solution driven by the installed base. Compaq's OpenVMS customers are satisfied with the outstanding capabilities of OpenVMS on AlphaServer systems and look to Compaq for server consolidation solutions when the capabilities of OpenVMS – e.g., clustering, availability, partitioning, and workload management – must continue to be made available.

Sidebar 1: Typical OpenVMS-Based Server Consolidation Business Drivers

- TCO reduction.
- Improved end-user service level/maintenance of service level agreements (SLAs).
- Floor space reduction.
- Consolidation of expensive datacenter real estate and redundant operations.
- Improved backup and recovery.
- Centralization of operations and reduction of duplicate efforts after a merger or acquisition.
- Improved server/storage use and performance.
- Platform and software standardization.
- Upgrade to newer hardware technologies to retire older systems with expensive maintenance contracts or inadequate performance or availability;
- The need for increased speed and flexibility in application development and deployment;
- Upgrade to newer hardware to replace "maxed-out" systems.
- Upgrade to newer operating system technologies to gain business benefits for both end-users and IT departments.

The sections that follow provide an overview of the basic technologies supporting OpenVMS-based server consolidation initiatives. Subsequent sections outline the capabilities and differentiators of certain important features of the Compaq OpenVMS operating system, OpenVMS Galaxy software, and OpenVMS Clusters. After this overview there is a discussion of several of the kinds of server consolidations that can be implemented, how they are accommodated, and three case studies. After reading this material the reader will fully understand the positioning of OpenVMS in server consolidation.

This report is part of a series covering OpenVMS Server Consolidation.[1] The business drivers for server consolidation and Compaq's methodology for such consolidation are discussed in detail in these additional documents.[2] The business drivers vary in type and importance according to the user environment. Typical drivers that apply to the OpenVMS environment are summarized in *Sidebar 1: Typical OpenVMS Server Consolidation Business Drivers*.

Many OpenVMS customers find themselves in a situation in which a number of VAX or older AlphaServer systems (of varying sizes) running multiple applications are combined into one large AlphaServer system or an AlphaServer OpenVMS Cluster.[3]

There are a variety of consolidation methods available. These methods include, for example, moving all applications to a more powerful machine (faster, more memory, more storage), or running on a single OpenVMS operating system instance on that large machine or a cluster. Given the raw power of AlphaServer systems, even the low-end AlphaServer DS Series may be the platform of choice.

1 The reports may be viewed by subscribers at www.dhbrown.com or at openvms.compaq.com/openvms/whitepapers for an even greater variety of information on related subjects.

2 *Enterprise Server Consolidation: OpenVMS Real World Server Consolidation Case Study and Enterprise Server Consolidation: Compaq Methodology, Issues and Best Practices*, D.H. Brown Associates, Inc., September 2001.

3 Compaq no longer provides new VAX machines either as stand-alone or clusters.

Sidebar 2: Compaq Server Consolidation Services

- For the consolidation planning phase of the consolidation there is the Consolidation Value Workshop and Server Consolidation Assessment Service.
- For the design phase, there is the Consolidation Design and Architecture Service.
- For the development phase, there is the Consolidation Implementation Service.
- For the actual migration there are Server Environment Upgrade/Migration Services, which include the GS Series AlphaServer QuickStart Service and the Application and the System Migration Service. The latter includes Migration Assessment, Application Migration Detailed Analysis and Design, and System Migration Detailed Analysis and Design.
- For the Support and Management phase there is the Outsourcing and Operations Management phase.
- There are high-availability services (including uptime guarantees).

In all of the above, the customer's individual needs determine the process entry point and the extent of involvement by Compaq Services.

For the high-end AlphaServer GS Series, it is possible to use Compaq's OpenVMS Galaxy software to set up multiple OpenVMS instances in partitions in either a single AlphaServer or AlphaServer cluster. The cluster architecture provides availability, scalability, and other advantages, and along with the Galaxy software provides major OpenVMS differentiators compared to other vendor offerings.[4]

As previously mentioned, depending on their current environment and future needs, IT departments may choose to implement a number of types and combinations of consolidation projects.

Physical (collocation) consolidation refers to the collocation of servers in a shared location. Benefits accrue in terms of economies of operational costs, such as system management, asset management and operations.

Server data integration (storage) consolidation involves standardization on a storage platform, which includes large shared RAID systems and backup solutions, as well as the integration of servers and SANs. Benefits accrue in terms of more efficient usage and management of disk resources, and better security.

Workload (application) consolidation refers to the collocation of applications onto fewer, more powerful servers. Benefits accrue in terms of reduced hardware and software expenses, as well as additional reductions in operational costs.

For the OpenVMS environment, the path to server consolidation focuses on and is led by services. Compaq's Services operation, along with the OpenVMS operating system, OpenVMS Galaxy software, OpenVMS Clustering capabilities, the AlphaServer hardware, and StorageWorks comprise Compaq's major differentiators in the OpenVMS environment. In partnership with the customer, a server consolidation solution is custom-crafted to specifically address one or more of the business drivers, and consolidation methods mentioned above. These services are discussed in detail in the documents referenced in Footnote 2 and are summarized in *Sidebar 2: Compaq Server Consolidation Services.*

[4] Discussion of the issues in determining whether a large AlphaServer or a cluster is required is beyond the scope of this document although certain aspects of clusters are discussed later.

COMPAQ OPENVMS ON ALPHASERVER SYSTEMS

Compaq's AlphaServer DS, ES, and GS Series systems running OpenVMS are an attractive server consolidation target for existing OpenVMS environments. This series of platforms provides high levels of performance, availability, and flexibility from the low-end through midrange and high-end server environments. The modular system architecture enables users to start with a system that meets immediate needs, but easily scales to larger systems as requirements change.

At the high end, AlphaServer GS Series servers are available in entry-level configurations with up to eight CPUs and two system partitions, and are expandable at the high-end to 32 CPUs and eight system partitions.[5]

The AlphaServer GS series is the latest high-end offering in the AlphaServer system family. The Alpha architecture has a track record of 10+ years of experience. Dozens of case studies[6] show that it provides a mature, very scalable, robust and stable 64-bit computing environment, key to the high-end demands of server consolidation.[7]

The sections that follow provide an overview of the capabilities of certain important features of the Compaq OpenVMS operating system, OpenVMS Galaxy software, and OpenVMS Clusters differentiators related to the needs of server consolidation. After this overview there is a discussion of the different types of server consolidations that can be accommodated, how they are accommodated, as well as three server consolidation case studies. After reading these sections, the reader will fully understand the positioning of OpenVMS in server consolidation.

WHY OPENVMS FOR AVAILABILITY IN A CONSOLIDATED ENVIRONMENT?

The OpenVMS operating system is well positioned to support large-scale mission critical applications in a consolidated environment. These features have been refined through many years of production use in real-world environments.

OpenVMS, for example, allows alternate pathing with automatic failover (transparent to applications) for storage, network peripherals, and cluster interconnects with a single server. In a cluster environment, if a node fails unexpectedly, OpenVMS provides direct access to system resources (CPU, memory, disk) through other cluster nodes in the cluster. This access means that

5 Further details on the AlphaServer hardware are available at compaq.com/products/servers.

6 See compaq.com/solutions/serverconsolidation.

7 Over a multi-year period starting in 2001, Compaq will standardize its 64-bit enterprise server product lines on the Intel Itanium Processor Family. Compaq announced that there will be one more follow-on high-end Alpha system. Compaq also promised to bring all of its operating systems, middleware and application portfolios to the Itanium architecture. This will be a smooth transition, the plans for which are discussed in openvms.compaq.com/openvms/roadmap/openvms_roadmaps.html.

the only reconnect required would be the session connection itself, i.e., simply re-establish the connection on another node in the cluster and the process continues. There is no failover of disk resources, application restart, or failover scripts required to be written, because an OpenVMS Cluster can run applications on all systems at the same time.

With OpenVMS Clusters it is possible to install dynamic system patches while applications are running and available to users in a multi-node cluster. These patches can be installed remotely so that the system network capabilities are not disabled during the patch operation. Because OpenVMS supports file multi-versioning, patches do not overwrite files in use. And, applications running on the target system operate while patches are being installed. Finally, patches can be backed out and/or removed.

OpenVMS also accommodates memory error resiliency through use of a just-in-time scrubbing model. Correctable errors are scrubbed when encountered by the operating system or application code. This capability ensures maximum availability. To enhance this capability, OpenVMS proactively locates and scrubs correctable memory errors prior to the operating system or application reading the memory. The memory page containing the error is marked bad and if the page is free (or when it becomes free), the page is mapped out so it will not be reused.

On designated high-end AlphaServer systems, OpenVMS also supports the online removal and replacement of failing CPU components.

One of the reasons why OpenVMS continues to be the choice for many mission critical environments is that it has long provided industry-leading disaster recovery capabilities with its unique active-active, load-balanced multi-site cluster technologies.

An active-active cluster not only shares the full read-write load between all servers at all sites, but it can also load balance the users – i.e., they are directed to the least busy servers in the nearest site available to them. Because this happens at the operating system level, no change in any cluster aware applications is required.

This provides a competitive advantage for OpenVMS. The shared cluster file system, distributed lock manager, and host-based volume shadowing features underlie this capability. These features allow customers to run the same applications in full read-write mode from all servers at all sites at all times, so the data is always 100% consistent on all sites.

Customers do not need to worry about data that might be lost during a replication window (the time between replication updates) and can maximize the use of the equipment at all the sites by sharing the full application workloads between sites (i.e., equipment avoids sitting idle at a backup site waiting for a primary site disaster).

This type of cluster ensures that data is 100% consistent at all sites at all times by using host-based shadowing (mirroring) software technologies to create "virtual" drives. From an application perspective, the system does not know (or care) that the "virtual" drive it is writing to actually consists of a physical drive (or hardware RAID partition for even higher availability) for each site.

STILL MORE FROM OPENVMS

As part of the base operating system's licensing (no additional charge), OpenVMS includes Compaq's Reliable Transaction Router (RTR). RTR offers a fault-tolerant middleware software product with transaction integrity. It provides Compaq customers and partners with a software fault-tolerant infrastructure that includes wide area disaster tolerance, continuous application availability, and ACID transaction integrity on OpenVMS (and other platforms).

RTR provides protection against multiple failure types in the distributed environment. These failure types include network, application, database, and system and site failures. When network, system, or site resources become unavailable, RTR automatically performs failover and recovery transparently to the application. This capability allows real-time, continuous service. Because of these features, RTR applications typically address stock and option exchanges, Internet banking and brokerage, electronic payment systems, trading and risk management, reservation systems, and telecommunications systems.

Compaq's Secure Web Server (CSWS) is another part of the base OpenVMS licensing (no additional charge). This part of OpenVMS is based on the popular Apache web (HTTP) server and allows:

- Use of full 64-bit AlphaServer hardware and operating system software for the very large scaling associated with centralized server consolidation initiatives.
- Full cluster and Single System Image (SSI) support. Full cluster and SSI support means it is easier to manage many web servers. For example, files common to the many web servers only need to be updated once in apache$common. Files specific to each server can still be uniquely modified in apache$specific directories.
- Non-susceptibility to the many 32-bit X86 server viruses common today.
- Customers can use the many open source programs based on Apache and at the same time remain confident that their web server stands on a secure operating system model. In view of the many hacker attacks that are common today, this provides added protection.

Furthermore, as part of the base operating system licensing, OpenVMS includes the Enterprise Directory for eBusiness. This enterprise directory product can be used to implement a distributed network directory service based on industry standard CCITT X.500 software that fully supports access via LDAP v3/v2 protocols.

Using the OpenVMS Enterprise Directory X.500 industry standard model, departments and organizations may adopt an incremental, independent approach to the establishment of a directory service using conforming products from multiple vendors. These separate implementations may then be connected to provide a single logical directory service that spans the department, organization, region, or the "world." The directory may contain information on people, systems, network resources, authentication certificates, and databases. Individual users and applications may access the directory.

Finally, for customers looking to migrate to newer software technologies as part of their consolidation planning, the base OpenVMS licensing (again at no additional charge) includes the latest versions for industry standard Java, Java servers, servlets, and Enterprise Java Beans (EJBs).

A GALAXY VIEW

Server consolidation based on the OpenVMS operating system enjoys the benefits of the latest software technologies to increase the utility of OpenVMS systems. OpenVMS Galaxy software, a major differentiator for consolidated servers running OpenVMS, for example, allows system administrators to make much more efficient use of existing system resources. Most importantly, existing applications do not have to be rewritten to use Galaxy, which is managed through the AlphaServer Console.

Licensed separately from the OpenVMS license, OpenVMS Galaxy software works with the AlphaServer GS Series systems to efficiently cope with changing computing demands. OpenVMS Galaxy flexibility allows the benefits of server consolidation to play an even larger role and helps server consolidation to reduce cost of ownership even further than might otherwise be possible. Why is this the case?

With OpenVMS Galaxy, it is possible to:

- Run multiple instances of the OpenVMS operating system (complete operating system copies tunable for specific applications with each instance using private and shared memory) on a single AlphaServer system in a single (or multiple) hard partition of the high-end AlphaServer GS-Series (GS80, GS160, GS320). OpenVMS Galaxy is itself an implementation of soft partitioning. (See *Sidebar 3: Hard and Soft Partitioning*).
- Enhance scalability since each instance can use an SMP architecture and the instances are clustered using a shared memory interconnect. This capability bypasses both the traditional SMP bottleneck and the distributed lock manager latency bottleneck.[8]
- Assign CPUs and I/O to each operating system instance dynamically as needed to eliminate bottlenecks.

[8] Further material on this bottlenecking bypass may be found at openvms.compaq.com/availability/galaxy.html.

- Dynamically allocate CPUs on the fly (without rebooting or disturbing running applications) as workloads change (using drag-and-drop or automatic load balancing).
- Cluster (with up to 96 processors) multiple operating system instances within an AlphaServer system or with other OpenVMS systems (within a single box or between multiple boxes) to provide high availability and disaster tolerance (up to 500 miles [800 kilometers] with an asynchronous connection).

Sidebar 3: Hard and Soft Partitioning

The AlphaServer GS Series offers both soft and hard partitioning. A hard partition is a system resources subset that runs an operating system instance. To varying degrees, depending on the implementation and other factors) hard partitions are isolated from each other electrically so a failure in one partition does not cause a failure in another and no resource contention occurs.

A soft partition consists of a system resource subset. However, partition isolation is less stringent. As a result, issues in one soft partition can affect another soft partition. Hard partitions are often broken up into soft partitions. Note also that partitions can be static and require a reboot to establish the new resource configuration, or they may be dynamic, and require no re-boot upon reconfiguration.

When using OpenVMS Galaxy, the OpenVMS operating system controls the access from one soft partition to another so that resources can be dynamically assigned from one to the other. Multiple Galaxy-based soft partitions can exist within an AlphaServer GS-Series hard partition. For example, a GS-Series 80, 160, or 320 AlphaServer can be configured into two hard partitions. What is more, one hard partition can for example, run Tru64 UNIX while two instances of OpenVMS can run in two soft partitions of the second hard partition. This capability of running multiple operating systems on the same hardware, and even different versions of the same operating system, offers a major benefit to anyone concerned with enterprise server consolidation.

OpenVMS Galaxy is managed through the Galaxy Configurator Utility (GCU). This GUI-based tool allows system managers to combine both Galaxy partitioning features with workload management.

The GCU, for example, allows system managers to define how CPU resources among multiple operating system instances will be shared (e.g., by time of day, by performance rules, or manually). When changes are made to sharing policies, no corresponding changes need be made to underlying applications or databases. This aspect is a key concept of GCU.

The OpenVMS Galaxy capabilities affect cost of ownership in a most positive manner. For example, through the use of Galaxy software, existing computers can be upgraded, and their capacity expanded. Savings in maintenance, system management, floor space usage, power requirements, and software licenses mount rapidly – especially because today's IT corporate directive to IT managers calls for doing more with fewer resources. These cost of ownership savings due to OpenVMS Galaxy and other OpenVMS capabilities have been well documented.[9]

OPENVMS CLUSTERS

The OpenVMS Cluster, another Compaq OpenVMS differentiator, puts forward a highly integrated combination of software, VAX and/or AlphaServer computers, storage devices, and other shared peripherals (See *Sidebar 4: Why a Cluster?*). The cluster can include all computer types – from desktop-capable to datacenter-capable. In the OpenVMS environment, all can work together as a single, virtual system sharing files, resources, storage, queues, and more.

[9] See openvms.compaq.com/openvms/whitepapers.

OpenVMS (on the high-end GS-80, 160, and 320 Series) with the aid of the Galaxy software (discussed above) supports one to eight dynamic hardware and/or software partitions with support for heterogeneous operating systems and operating system versions. These machines with OpenVMS Galaxy can also be clustered.

Advanced features include:

- fully shared, multi-node, read-write disk access;
- a cluster-wide file system;
- distributed lock manager;
- cluster-wide batch/print subsystem;
- votes/quorum-based membership management;
- shared system disk;
- single security domain; and
- single system management domain.

Beyond this, there is also cluster-wide application programming interfaces, APIs, capability for mixed VAX and AlphaServer nodes, rolling upgrade support, and multiple interconnects (CI, Memory Channel, Fibre Channel, SCSI, Ultra SCSI, DSSI, LAN, WAN). Support is provided for up to 96 clustered nodes, built in failover and load balancing, cluster network alias, disk/tape servers, and the multi-site disaster tolerance capability discussed earlier.

Many of the earlier features alone offer justification for choosing OpenVMS clusters (e.g., cluster-wide file system, 96 node capability spanning up to 500 miles, etc.). Following September 11, disaster tolerance, high availability, and security have taken on heightened significance.

OPENVMS CLUSTER AVAILABILITY

OpenVMS clusters reduce both unscheduled and scheduled downtime, provide fast recovery time for hardware and software errors and superior recovery procedures, balance workloads across servers, and simplify resource management. The clusters operate 24 hours a day, 365 days a year, and they can run at 99.999% uptime with an optional guarantee. An additional reason for this uptime is the OpenVMS cluster's ease of management. (See *Sidebar 5: Easier Management.*)

When most users talk about high availability, they claim to need "many nines" availability. Unfortunately, few vendors include scheduled downtime for proactive maintenance such as operating system upgrades, tuning reboots, hardware upgrades/replacements, and the addition of memory, I/O controllers, and disks in their many-nines availability numbers.

Failure to include scheduled downtime in a system availability number is unacceptable in today's mission critical environments. For example, today's environments include more users than ever before. This is even more of an issue

for server consolidation environments because the number of users affected is much higher than with traditional servers. It is difficult to tell so many users that they are being disconnected in order to failover a server.

The bottom line is that users want to get as close as possible to 100% application availability, where that availability includes scheduled, as well as unscheduled, system downtime.

Users in multiple markets often need to be able to measure their application uptime in years, rather than in hours per year. OpenVMS clusters satisfy this requirement. For example, Irish National Railway shut down its OpenVMS Cluster-based application after 17 years, and then only because they moved to a process control-type system.

An important availability characteristic in clusters is separating system availability from application availability. This OpenVMS clusters capability allows system managers to shut down individual cluster servers for scheduled maintenance with no affect on application availability. The significance relates to how users view availability. Users do not care if system managers reboot servers as long as their application remains up and running and their system view shows that they are always connected. (This feature is explored in more depth later in this report.)

Sidebar 4: Why a Cluster?

When an organization chooses server consolidation, platform choice remains an issue. A server consolidation project requires selection of one or more platforms on which to consolidate applications and databases that currently reside on many servers.

Analysis shows that some user companies have actually moved from 300 servers before the consolidation to 15 servers afterwards and even fewer if the companies moved to a clustered platform. The savings is, of course, environment-dependent. Such consolidation means that the consolidated servers may host 500 to 1,000 users per server, rather than the more common 50-150 users per server in an unconsolidated server environment. This large number of users makes availability, reliability, and scalability major requirements for a server consolidation platform. Clusters perform very well in these areas.

The platform choices for a consolidation range from server clusters to large SMP machines and mainframes. Single SMP servers often save costs in consolidating workloads, but they may not provide the desired availability. Multiple architectures are evolving to partition SMP platforms to isolate hardware and software resource groups from each other. With these partitions, should a failure in one partition occur, resource groups and applications in other partitions are not affected. This is what OpenVMS Galaxy features provide.

Such partitioning has been available on mainframes for many years, and SMP manufacturers are developing similar partitioning schemes for their SMP servers. Since availability remains a major reason for server consolidation, the case can sometimes be made that single SMP systems alone are not suitable for server consolidations. Clusters are deemed better, based on the cluster advantages. This, of course, depends on the customer's needs as determined by a careful analysis.[10]

A cluster offers an integrated combination of servers (sometimes but not always more modestly sized than large SMP servers) designed to achieve continued operation and availability of data and applications after failure. Such availability cannot be obtained with non-clustered machines except through very expensive fault-tolerant machines.

Note: Fault-tolerant solutions depend on active-passive (site replication) multi-site solutions. With its multi-site, active-active, load-balanced capabilities, OpenVMS Clusters also provide a degree of availability not otherwise generally obtainable – even with fault-tolerant solutions.

[10] Compaq Services can assist in the decision-making process.

ALWAYS CONNECTED

OpenVMS handles the just-mentioned 100% availability during scheduled downtime in the following manner. Consider a three-node OpenVMS cluster. When system administrators need to shutdown a server for maintenance, they set flags and/or disable new logins on the server to be shut down. This technique allows all current database connections and sessions to continue on that server. However, it diverts all new connections to the other cluster servers.

After a time (environment-dependent), system managers shut down the server scheduled for maintenance with zero impact on application availability. There is no need to failover any resources. There is no need to notify users because they are unaffected. And after maintenance, when the server is brought back online, logins will be enabled and the server will automatically be part of the cluster's load balancing system.

This environment is what high-availability users really need now, and will need even more in the future as downtime becomes even less acceptable due to the ever-increasing cost of lost business. Users need this kind of availability because upgrades, tuning reboots, and hardware updates are maintenance tasks that must be done on any platform. For example, while some vendors, like Compaq and its AlphaServer GS 80/160/320's (running OpenVMS Version 7.3) support CPU hot swapping, it is not yet possible to hot swap memory. To add new memory to an OpenVMS cluster, however, system managers follow the procedure described above to shut down the server for maintenance in a manner that does not affect application availability.

Should a failure occur in a server that is part of a multi-site cluster with full read-write and synchronized data capabilities, as well as load balancing capabilities, it has little significance, because applications do not even know which server in what datacenter processes their requests. In the event of a datacenter catastrophe, the business can even continue without any data loss. Application availability is not disrupted when the customer has installed a disaster tolerant configuration. Note that this is particularly important since the events of September 11.

In the event of a system failure in an OpenVMS cluster, only the application or database connections that resided on the failed server need to reconnect to another, least-busy server and continue processing. All application or database connections that are being used on other servers do not have to reconnect.

No system failover, application restarting, or disk remounting is required because the application and drives are already mounted and running on the other servers. Application processing merely continues.

Sidebar 5: Easier Management

It is typically more difficult to manage a cluster, which is a collection of servers, than to manage a single large SMP machine. OpenVMS simplifies cluster system management because OpenVMS's Common System Disk (i.e., single-system image) allows system managers to oversee an entire cluster as if it were a single server.

With a Common System Disk, each server boots up only one system disk image for the entire cluster. System managers can see and manage the entire cluster as a single server – that is, one user database, a one-batch/print queue subsystem, etc. The root file system on each node in the cluster is part of the cluster file system. All the cluster servers can directly access all cluster storage for read-write I/O. There is no requirement for a server to "serve" storage to other servers, unless the system administrators want such a configuration.

Another advantage of a single-system image is the ability to upgrade not only applications, but the operating system for all the servers and users at once. For example, if four servers run off a Common System Disk, it is necessary to upgrade the system disk only once for all users to run a new operating system version.

The disadvantage of using the single-system image when upgrading the operating system (not the applications) is availability disruption because the administrator must disconnect the users to perform the upgrade at one time. To address this issue in an OpenVMS cluster, system managers generate two single-system images (or two Common System Disks), but share common cluster-wide user database queues etc. and other files between them.

When the system managers are ready to perform the operating system upgrade, they migrate the users from one Common System Disk environment to the other. When the system comes back up after the upgrade, the users are failed over to the environment using the Common System Disk that was just upgraded. Then the other Common System Disk environment is upgraded.

Batch queues can run jobs on any cluster node because the cluster has a cluster-wide batch system. This means the cluster maintains a generic batch queue, which feeds specific queues on specific servers. The system manager can regulate the number of queues, the streams per queue, and the base priorities per queue.

Because every server can see all the storage, individual jobs need not know which server they are running on. With this system, when a user submits a job to be run, that job is automatically submitted to the generic batch queue. Then the batch queue determines the least busy system on which to run that particular job.

In the event of a server failure, OpenVMS clusters experience a brief delay, typically less than 5-10 seconds (but certainly dependent on the applications, database, number of users, etc.) in processing while resources are re-allocated. But, the overall application remains available and only the connections that were open on that one failed server need to reconnect to another available server. The amount of time required for a cluster to continue processing from the same place is minimal.

STILL MORE OPENVMS CLUSTER TECHNOLOGY

The OpenVMS Cluster provides a number of advanced technologies for Compaq partner and user benefit. Consider the Distributed Lock Manager (DLM), for example, which is at the heart of an OpenVMS Cluster. The DLM is used extensively by the operating system, and is available for application use. The DLM typically, but not necessarily, coordinates all file operations. It provides asynchronous completion, asynchronous blocking notification, deadlock detection, lock trees, lock quotas, dynamic remastering, value blocks, and more. Other technology highlights in OpenVMS Clusters include,

- Load balancing at many levels. The OpenVMS Cluster I/O subsystem, the DLM, the batch/print subsystem, cluster communication traffic, user logins, and network links all can be configured for optimum load balancing.
- To deal with scaling needs, entire systems and I/O controllers can be added dynamically to an OpenVMS cluster with zero impact on application availability. Storage subsystems can scale on multiple interconnects to many terabytes of online data.
- Single security Domain. OpenVMS Clusters implement exactly the same security mechanisms as single OpenVMS systems. All access controls, policies, alerts, auditing, logging, and security management functions operate cluster-wide as if the cluster were one system.
- Single disk quotas across all servers. It does not matter what server in the cluster the user connects to, the disk quota assigned and monitored will be the same.

Batch queues are another advantage. Batch queues can run jobs on any cluster node because the cluster has a cluster-wide batch system. OpenVMS Clusters provide the concept of generic queues and execution queues. The system manager can regulate the number of queues, the streams per queue, and the base priorities per queue.

Jobs can be submitted to generic queues that feed the job to any of several execution queues on different servers. For example, if a job must run every day at midnight, it can be submitted to a generic queue.

Every evening at midnight, the batch subsystem identifies an available execution queue on the least busy node on which to run the job. From an operations and end user perspective, the advantages of a cluster-wide batch queue system derives from a batch job being written without worrying what node or server it runs on. All servers have exactly the same view of available disk resources.

The Cluster-Wide Process Control API is another advantage. OpenVMS Clusters provide an API for cluster-wide operations. Through use of this API, a process can create, monitor, control, and delete another process on any cluster node. Processes can elect to be notified of cluster membership changes. System utilities take advantage of these abilities and can display information for the entire cluster, rather than only for the node on which the utility is executed.

All OpenVMS nodes in an OpenVMS Cluster can issue direct I/O requests to storage devices at the same time. Routing I/O activity to a device through a single node is not necessary (the usual case with shared nothing cluster architectures). The DLM automatically coordinates access to shared devices. This feature greatly increases the I/O throughput.

In addition, with appropriate network products, OpenVMS Clusters provide a network alias feature that allows incoming network links, either TCP/IP or DECnet, to be targeted to the cluster by a single name. The load balancing

software chooses which actual cluster node (typically the least busy) services the network link. Outgoing network traffic can also use this feature, so that, for example, e-mail leaving the cluster uses the cluster alias name rather than the name of a specific node.

Finally, there is Rolling Upgrade Support. OpenVMS Clusters configured with more than one system disk can run different OpenVMS versions on each system disk. This capability allows OpenVMS Cluster nodes to be upgraded in a rolling manner. This process allows upgrades to be done over a period of time to take advantage of new operating system and layered product features. In suitably configured clusters with multiple system disks, these operating system upgrades can also be done transparently and with zero application availability impact to the end-user community.

CONTINUOUS COMPUTING AND APPLICATION AVAILABILITY

All of the above-described technology plays well together in today's computing environment, which is much more difficult to manage than in the past. For example, with Internet traffic loads, e-business applications, new technologies (LDAP, XML, J2EE, X.509 security certificates) and much tighter B2B integration between vendors, customers, and end users, it has become more difficult for IT operations groups to schedule planned system downtime. Such scheduling usually means affecting application availability in a negative manner.

As noted above, it is not acceptable in many cases to tell partners, end users, and customers (perhaps spread across the globe and/or Internet) that they will be disconnected to shut down a server. Even if a backup server is available, it often means disconnecting current users and forcing them to reconnect to the backup server.

Separating system availability from application availability (OpenVMS Continuous Computing) offers an important and unique availability capability with OpenVMS Clusters. This OpenVMS capability allows system managers to shut down individual cluster servers for scheduled maintenance with zero application availability effect. Such a capability is important because business groups usually do not care if system managers reboot servers as long as their application view shows that the servers are 100% available (no disconnected users or application restarts). This concept calls for a further drill-down.

OpenVMS allows 100% application availability during scheduled/planned system downtimes. Consider, for example, a three-node OpenVMS cluster that has been configured such that any two nodes can handle peak loads for short periods of time. When system managers need to shutdown a server for maintenance, they set flags and/or disable new logins on the server to be shut down. This technique allows all current database, user, and application connections to continue until completion.

All new connections are diverted to the other cluster servers. Any new batch jobs are started on other servers in the cluster. Because OpenVMS provides a shared file system with its DLM, there is no need to fail over any served disk resources or restart any applications, as is often the case with other operating system platforms. These disk resources are already available on other OpenVMS servers. New users are simply connected to other servers transparently.

As discussed earlier, after an environment-dependent once all open application and/or database connections are completed and current batch jobs are finished, the system manager can shut down the server with zero impact on application availability. Notifying users that a server is going down need not occur because they will not be affected. Once the maintenance or hardware upgrade is completed, when the server is brought back online, logins are enabled and the server will automatically begin participating again in the cluster's load balancing scheme.

For a multi-site, load-balanced cluster, if there is a datacenter outage, the business can still continue without any data loss at all as shadowed data is always 100% consistent at all sites, at all times. When users or applications reconnect, they will be directed to the servers in the remaining datacenter.

Such an environment is what high-availability and continuous computing users require today. They need this kind of availability because operating system upgrades, tuning reboots, and hardware updates to accommodate rapidly changing business requirements are tasks that must occasionally be done on all platforms. In the future, due to the ever-increasing cost of lost business, users will start to demand even more of these types of capabilities as application downtime becomes less acceptable.

Using the background provided above about the numerous Compaq OpenVMS differentiators, the following sections provide an overview of the types of server consolidation that can be accommodated.

HOMOGENEOUS WORKLOAD CONSOLIDATION

Within the AlphaServer environment, one basic consolidation scenario consists of consolidating the homogeneous workloads of some number of smaller and/or older VAX hardware and/or AlphaServer systems into a DS Series (low-end), ES Series (midrange) or GS Series (high-end) server. The chosen platform is a determined by a sizing analysis that can be done either by the customer or Compaq Services.

The consolidation method is straightforward. Applications and data are moved to the new, larger system via a backup/restore, or by moving the disk drives to the new system (if the disk drives are to be kept). Each workload is then brought up on the new system. In fact, this process predates the term server consolidation, and has been accomplished in various environments for many years.

The key technology here is the availability of faster, higher-capacity systems on which to consolidate. The OpenVMS operating system easily supports large numbers of processes on high-end systems. And there is no need for the homogeneous applications to concern themselves with unpredictable interference between them. OpenVMS also provides a class scheduler, which can be used to ensure one application process does not exceed the CPU resources it is allocated.

Upward scaling capability is established by the power of the largest available GS Series server. For compute-bound workloads, the GS Series provides the availability of faster processors, which as well as up to 32-way SMP allows individual workloads to run faster, with more concurrent workloads supported. I/O-bound workloads benefit from more available bandwidth, with up to 244 PCI adapters connected to a single AlphaServer GS320 system.

Due to the availability characteristics of the GS Series servers, there is little concern for increased application failure due to scaling. The AlphaServer GS Series supports the replacement or upgrade of critical components without any disruption to users. Other reliability and availability features of the GS Series are summarized in *Sidebar 6: AlphaServer GS Series Availability Features and Functions.*

Customers with less performance or size requirements should investigate the capabilities of the midrange AlphaServer ES Series. Also not to be ignored is the AlphaServer DS Series. The AlphaServer DS system is ideal for state-of-the-art, low-end performance requirements and can often replace numbers of older, lower performance VAX machines.

Sidebar 6: AlphaServer GS Series Availability Features and Functions

- Online Maintenance and Upgrades – Most of the AlphaServer GS Series components are designed to support online maintenance and upgrades. Depending on their function, components can be hot swapped (removed and inserted without shutting down the system). OpenVMS supports operating system, database, and applications upgrades.
- Hardware Monitoring and Environmental Reporting – System components and modules and the operating environment are monitored, with changes reported automatically as they occur for appropriate action.
- Fully Redundant Service Processors with Automatic Failover Capability – System components and modules have a dedicated monitoring processor. This network of monitoring processors is constantly available, surviving component failure or the absence of main power. Tracked information includes model variants, serial numbers, firmware levels, and component status, including loss of power. Changes are reported to the system console. The console is not a single point of failure. A console failure results in automatic and transparent failover to another console.
- N+1 Fault-Tolerant Power Subsystems – N+1 fault-tolerant power capabilities offer automatic and immediate failover, with no disruption to applications if a power supply fails. Optional dual AC inputs to separate AC power grids are available for those concerned with external main power disruption.
- High-Availability Networking and Storage Subsystems – AlphaServer GS Series servers can be connected to highly available, redundant external storage and network subsystems to comprise a highly available system configuration. Input and output multi-pathing is supported for both storage and networks, to allow a high degree of data and transactional integrity.
- Easy Serviceability and High Availability – The AlphaServer GS Series benefits from Compaq's extensive analysis of causes of component and system failures. Serviceability and availability features include fewer components and cables, point-to-point links instead of busses, light-emitting diodes in CPUs for easy identification, and color-coded building blocks. In addition, cables and modules make identification easy for service engineers, self-aligning connectors guarantee a perfect connection without the risks of bent pins or off-center contacts, and all serviceable hardware uses captive fasteners. Moreover, only the highest quality components and modules are used in critical subsystems. Memory, cache, and all system data and instruction pathways are protected with ECC (error correcting code) logic to prevent data corruption and resulting system failures. Finally, extensive factory burn-in avoids "infant mortality" problems.

HETEROGENEOUS WORKLOAD CONSOLIDATION

Assume that it is necessary to consolidate the workloads of smaller VAX and/or AlphaServers onto a larger AlphaServer. Further assume an additional issue is that the workloads are heterogeneous, and cannot coexist on a single operating system instance.

Each smaller machine is replaced with a separate OpenVMS instance running in its own partition on a GS Series system. The operating system, middleware, and application software are moved, unchanged, to the new system and rebooted. As in the previous homogeneous application consolidation scenario, disks can be moved unchanged, or the data can be loaded in any convenient way onto new storage.

The key technology brought to bear here is partitioning. There are two kinds of partitioning. Hard or system partitioning (hardware partitioning); and application/resource partitioning, also referred to as software partitioning or resource management. Hardware partitions occur on AlphaServer quad building block (QBB) boundaries, thereby allowing up to eight partitions on a 32-

processor system.[11] With hardware partitioning, it is possible to run different versions of OpenVMS (or even other operating systems or operating system versions) in different partitions.

An application or operating system failure in one independent system partition has no effect on applications running in other partitions. This capability can be of great practical utility, especially in environments where individual business resources must be kept separate, or where certain applications that are not well behaved must reside on separate systems. For better-behaving workloads, the use of OpenVMS resource management tools within a single partition eliminates the need to reconfigure the partition (while still managing the workload).

RESOURCE MANAGEMENT TOOLS

A brief discussion of some of the tools available to further increase the value of multiple instances of OpenVMS and partitioning is in order. For example, resource management tools complement system partitions by providing more fine-grained system control.

OpenVMS offers a variety of resource management capabilities within the operating system. These capabilities provide system administrators the flexibility to tailor a system to fit specific business requirements, as well as to predict and guarantee service levels. The resource management capabilities make it possible to allocate, control, and monitor CPU and memory resources for individual users, groups of users, or specific applications within a single OpenVMS operating system instance.

Such capability is accomplished by applying a resource consumption policy that directly supports the business goals of an enterprise. This policy is dynamically enforced at all times, so that in the event of unexpected load fluctuations, critical applications will not be resource-deprived.

OTHER TOOLS

There are still other tools of great utility. The Compaq Galaxy configuration Manager is an intuitive management GUI-based (graphical user interface) application that enables an administrator to respond to changing workload requirements by dynamically moving CPUs between different instances of the operating system.

OpenVMS supports a variety of tools to manage CPU resources. With processor sets, individual CPUs can be grouped and dedicated to specific applications or users. CPUs can then be moved between processor sets administratively or, should one set require additional resources during peak load periods, through the resource usage monitoring capabilities provided in OpenVMS.

[11] On an AlphaServer GS server, the smallest hardware partition on a QBB consists of four Alpha microprocessors, four memory modules supporting 32 GB of memory, and eight PCI buses with 28 PCI slots.

Class scheduling adds further definition to how CPU resources can be managed by assigning users and applications to a class, which is allotted a percentage of the CPU time available on the system (as a whole or within processor sets). Many different classes may be defined for a given system or system partition. This definition capability enables an administrator to apportion the necessary CPU cycles among all the system jobs. This effectively eliminates the risk that a highly compute-intensive job will dominate the system and degrade services for other users.

MANAGEMENT/OPERATIONAL CONSOLIDATION

In this scenario, existing AlphaServer GS Series systems, each with its own management and operational domains, are consolidated into one management/operational domain. In addition to the savings in system management, the availability features of the GS Series platform can be further leveraged.

This consolidation scenario is accomplished by establishing (or expanding) a cluster of systems. An OpenVMS cluster consists of from one to 96 nodes and spans up to 500 miles (800 kilometers). The cluster software runs on one or more AlphaServer systems connected together via a high-bandwidth low-latency Memory Channel or a LAN-based network interconnect.

The single-system image and shared file system capabilities of the OpenVMS Cluster greatly simplify common operational activities, such as application maintenance, patching, and account maintenance. Changes only need to be applied once, and apply to the entire cluster. A common event log makes system monitoring more efficient by logging all events to a common reporting facility.

OpenVMS Cluster technology also provides enhanced availability important to server consolidation environments, automated load balancing, single-system management, the Cluster Application Availability framework, and scalability for the entire cluster-wide file system.

Clients connecting to an OpenVMS Cluster system are assigned to the cluster member (either standalone server or partition) with the lowest current load. This continuous balancing of client connections provides optimized performance and resource usage.

As noted earlier, an OpenVMS Cluster system implements cluster management as a single system rather than as a collection of multiple systems. This implementation substantially reduces management and operational costs. The cluster-wide file system also contributes to improved data management, higher availability, and faster and easier application failover and lower costs by allowing cluster-wide storage management.

In the AlphaServer GS Series cluster environment, high-end scaling is accomplished by a combination of Scale Up (SMP machines) and Scale Out (clustering), in various combinations.[12] On a GS series system, a single OpenVMS instance can span from one or two partitions of four or eight CPUs in a GS80 system, up to eight partitions on a full 32-way GS320. A total of up to eight OpenVMS instances can be clustered together on a single AlphaServer GS320 system.

The combination of partitioning and clustering provides the flexibility to expand or reconfigure systems to meet nearly any performance or availability need, by creating or consolidating partitions within a single server, as well as expanding the cluster to incorporate additional servers and/or partitions. Simply put, the most stringent requirements can be met.

APPLICATION CONSOLIDATION

A common cause of server sprawl is the proliferation of small servers to handle file and print services for local workgroups. In this scenario, an appropriately sized GS Series server (or cluster of servers) running OpenVMS Advanced Server can provide file and print services (as well as other services).[13]

OpenVMS file and print services are well known and have been established for many years. Differentiators include,

- Not susceptible to the many 32-bit server viruses that are common today.
- Full load balancing across all cluster members.
- Additional servers can be added dynamically to the cluster with no downtime of any server. Cluster support is up to 96 nodes. It does not use a two-node cluster primary-backup cluster, as do other solutions on the market today.
- Can participate as a primary or backup domain controller or member server in a Windows NT4 environment or as a member server in a Windows 2000 environment. Domain controllers can be set up to transparently failover to other cluster members, with no user impact.
- OpenVMS is a true 64-bit solution.

[12] An ES-Series system does not Scale Up as high as a GS Series system. Otherwise, the discussion in this scenario is applicable to this hardware.

[13] Those who do not need the GS Series capacity can take advantage of this scenario using the ES Series or the DS Series.

ORACLE DATABASE CONSOLIDATION

An environment may have multiple Oracle databases, and/or multiple instances of an Oracle database running and want to consolidate them onto a large AlphaServer GS Series[14] database server. This need may be driven by Oracle pricing schedules. Oracle currently offers a processor-based pricing option that is particularly interesting for server consolidation.

How does all of the previous technical and business analysis in this report play out in the real world? The following three case studies show that realistic, functional server consolidations using the technologies discussed above may be found in different environments with substantial benefits in cost of ownership, availability, performance, and meet a variety of other business driver needs (see Sidebar 1). Still other case studies may be examined in a variety of sources.[15]

CONSOLIDATION CASE STUDY ONE – CANADIAN GOVERNMENT

Through server consolidation, one of the Departments in the Canadian Government improved the level-of-service provided to its users, made the systems more available and manageable, and made it easier for users to access the information they needed. How was this feat accomplished? The Department consolidated 41 heterogeneous standalone servers down to six regional clusters, all standardized on AlphaServer clusters.

The Canadian Federal Government Department server consolidation effort reduced weekly outages for scheduled downtime from four hours to one hour for older databases that had no built-in availability capabilities, and it improved availability in general for all involved systems. There was also a significant increase in ease of system management and security management because the consolidation made it necessary to monitor and manage only six environments instead of 41.

While cost of ownership was reduced according to those involved in the project, the dollar savings were not measured because the prime project goal was service-level improvement rather than cost savings. The consolidation also resulted in a number of not-easily quantifiable productivity benefits, as well as lifestyle benefits for the support staff who no longer had to spend weekends and nights on maintenance.

---

[14] Customers can take advantage of this scenario using ES Series servers if an evaluation of present and future Oracle workload requirements indicates that this midrange Series is sufficient. Although the ES Series server handles less Oracle workload on a per-processor basis, Oracle pricing is the same as the GS Series.

[15] See http://www.openvms.compaq.com/solutions/serverconsolidation/.

IT managers can drive a key message from this case study: Server consolidation goes beyond UNIX and/or Windows. It places emphasis on picking the right environment for the job at hand. In this case, the right environment was consolidation in the venerable, yet totally modern OpenVMS environment on AlphaServers.

WHAT WAS CONSOLIDATED?

The Canadian Government Department had an OpenVMS-based mission critical application distributed across one Headquarters site and five Regional Data Centers. Before consolidation, the application depended on a mix of 41 VAX and AlphaServer standalone servers. The application was distributed across all these machines via the DECNet protocol. The applications and users also used TELNET- and FTP TCP/IP- based applications.

Each of the 41 servers ran its system disk and its own copies of user products. The 41 servers had to be managed individually. Moreover, users had different login names on each server.

The main reason for choosing a centralized cluster solution was the organization's distributed computing environments, which made manageability and support difficult and time consuming. Also, through the server consolidation solution, the Canadian Federal Department wanted to make it easier for users to access information.

A main reason for choosing an OpenVMS Cluster as a server consolidation platform was the existing presence and familiarity with OpenVMS systems. The other reasons included the availability and scalability of OpenVMS Clusters, which makes them suitable as a server consolidation platform, as well as numerous, free resource management tools and server consolidation planning tools.

ISSUES THAT LED TO THE SERVER CONSOLIDATION SOLUTION

The Canadian Government agency found it difficult to maintain standard configurations throughout the organization. Over the years, the Department had purchased different devices and tools for the 41 separate servers. Unfortunately, the difficulties in maintaining standard configurations resulted in costly efforts to manage, license, monitor, repair, and upgrade the infrastructure.

In some cases, the organization used different products for different groups because it was too difficult to maintain a single standard across the servers. These non-standard products made application upgrades difficult and labor intensive because, in addition to the large number of systems, each had minor configuration differences.

Due to the nature of the application, all 41 servers had to be updated at one time on a weekend. This was difficult to manage between Friday night and Monday

morning. To compound the problem, the upgrade took the application offline for the weekend. The latter downtime generated still more problems.

The pre-consolidation environment labored under downtime problems due primarily to weekly database maintenance requirements. The problems occurred because the database that the organization used was old and had no availability built into it. The organization performed only weekly database import/export for performance reasons. If the server was down, users had to wait.

Users operating on AlphaServers obtained much better response than the VAX users. Consequently, one of the server consolidation goals was to give all users the same performance level. This was solved easily once all users ran on regional based clusters.

Because of the distributed environment, backups were performed over the network during off-hours. However, many users worked during off-hours. When backups ran, they caused network performance degradation. The OpenVMS clusters provided a means to do local backups of all cluster data, and thus removed this daily backup load from the network.

Naming differences were also a problem. Users had separate accounts on different class servers. The end users typically had different login names on the different OpenVMS and Windows NT servers. This issue required users to remember multiple login names and passwords. Also, users were given one print queue name in OpenVMS servers, and a different print queue name under Windows NT. These user account and print queue name differences were resolved during the consolidation.

AFTER CONSOLIDATION

"Extremely pleased" describes the reaction of the Canadian Government Department to the consolidation. Why did they have this reaction? After consolidation each geographic region had one application environment. There are three nodes in one cluster in each region. From the viewpoint of the IT and support staffs, the organization moved from having to upgrade 41 separate application environments on a weekend to having to upgrade only 6. The consolidation gave the support staff much more manageability and resulting scalability.

The organization's weekly outages for scheduled maintenance downtime on the older but critical databases were up to 16 hours per week (832 hours of downtime per year) on the smaller VAXs before the server consolidation. This is an unacceptable downtime in critical environments. The consolidation reduced the outages to one to four hours on all databases.

Overall, the Canadian Federal Department significantly increased its application availability, partly because of the cluster-based platforms. As an example, consider the following bad news and good news scenario in the Pacific Region cluster. The

bad news: A server failed at 4:00 AM because of a memory problem. The good news: By 11:00 A.M., nobody had noticed the failure. The reason the failure went unnoticed was the reconnection of the online users to one of the two remaining servers in the cluster, with no failover time, restart time, or initialization time, or need to remount disk drives.

For the users who had been running on the slower VAXes, the large Alpha Clusters appeared to run like the wind. The users previously running on AlphaServers also were delighted for two reasons. First, the new AlphaServers were faster than the older ones (as would be expected). Second, with the three-node cluster concept, each server was configured so that administrators could lose a node and still handle the cluster peak loads.

Overall IT support costs and complexities due to issues such as change, configuration, storage, security, and performance monitoring were reduced. The support cost reductions derived from a simplified environment because of the decreased numbers of server environments, as well as the single management view of the clusters.

As expected, the combination of the server consolidation and the basic features of OpenVMS Clusters decreased overtime costs for the support staff. In order to perform scheduled server maintenance in prime time with zero impact on application availability, the IT staff now uses OpenVMS's Cluster and TCP/IP Load Broker/DNS capabilities to route new users away from servers scheduled to be shut down.

Security has been increased as there are far fewer systems and accounts to monitor. As part of the increased security, the Canadian Federal Department found and rid itself of user IDs that were five years old and still on the servers. A consolidation process typically results in server cleanup.

Network traffic during off-hours has been greatly reduced because all cluster storage is now local to the cluster and all backups are performed within the cluster. Since the storage is local, none of the server backup traffic runs across the network. Consequently, users working on the network during the off-hours no longer have problems with backup data running at the same time.

Each region now has very scalable growth capacity. System managers can add CPUs within the box or they can add servers, memory, and devices as required. Moreover, they add hardware with no application availability impact.

User account management has been simplified in the consolidated environment. The Canadian Federal Government started planning a user account-naming scheme during the project's advance planning. At that time, the organization found that many users had one name for NT logins and another name for OpenVMS.

Because the organization had a single domain across Canada for MS Exchange e-mail, users already had a single unique NT name (equivalent to their e-mail name). Rather than changing all these NT/e-mail names, the server consolidation planning and implementation teams assigned the same OpenVMS user account names as those used for NT logins (which included access to the users' MS Exchange e-mail). These naming changes simplified server and e-mail access for users.

End-user printing naming also was simplified in the consolidated server environment. As was the case with user account names, the teams simply made all of the same print queue names on OpenVMS match the print queue names used for NT. Simplifying naming makes information access easier for users, and eliminates a cause of user errors.

Compaq Global Services played a role in this server consolidation. Throughout all the parts of the consolidation, throughout Canada, the services and consulting provided by this Compaq operation worked to ensure the cost-effectiveness and success of the complete consolidation endeavor.

CONSOLIDATION CASE STUDY TWO – VERIZON INFORMATION SERVICES

Dallas, Texas-based Verizon Information Services, a unit of Verizon Communications, is the world's largest publisher of print and electronic directory information. With approximately 20% of the worldwide market for directory advertising, Verizon Information Services provides sales, publishing, and other related services for some 2,500 directory titles in 47 states and the District of Columbia, and 18 countries outside the U.S. This effort includes more than 1,600 Verizon directory titles with a circulation of approximately 112 million copies domestically and 41 million internationally.

The Amdocs directory publishing system is the application behind Verizon's Yellow Pages. This system runs on Compaq's OpenVMS v7.2 and earlier OpenVMS versions. Verizon Information Services provides Amdocs customization so that the application meets customer requirements.

"The Amdocs application is key to Verizon's directory sales force," says Rick Hungerford, Advisory Systems Engineer at Verizon. It allows the sales force to sell advertising space and enter information into the system. It also documents sales and commissions, produces reports, and transfers the information necessary for publishing the yellow pages and bills advertisers.

As customer SLAs became more important, and cost control rose as a major issue, Verizon Information Services opted for both server upgrade and server consolidation to meets its growing needs. "Our VAX systems were maxed out,"

explains Hungerford. "We had a cluster of five VAX 7800 systems, the largest VAX systems available at the time, but we needed more processing power in order to meet our SLAs." While this analysis was underway, Verizon realized that consolidation would bring additional benefits.

Looking for virtually unlimited growth, they chose to upgrade to Compaq AlphaServer systems running OpenVMS. They were also well aware of the capability to run the Galaxy software in the future so they could invoke multiple instance OpenVMS partitions and move processing power about as needed. The three production AlphaServers in the new consolidated environment accommodate different geographic regions (three different time zones) and often need to serve peak load situations.

To get the horsepower it needed, and decrease total cost of ownership, Verizon Information Services replaced its production cluster of five VAX 7840 systems, each with 3.5 GB of memory, with three AlphaServer GS140 systems running OpenVMS 7.2, each with four processors and 14 GB of memory. In addition, there is a non-production cluster for development and training that consists of one AlphaServer GS140 and one DEC7700.[16] Both clusters have redundant controllers, power supplies, and power backup. Compaq Global Services takes care of all system maintenance.

As mentioned, a primary reason Verizon Information Services moved from VAX to Alpha systems was to take advantage of the extensive benefits of server consolidation and the firm was not disappointed. The VAX to Alpha upgrade and concurrent consolidation resulted in cost savings reaching 40% in software costs, 25% in maintenance costs and 20% less floor space. Moreover, it now takes six people to maintain the systems rather than eight.

More benefits accrued. For example, Verizon Information Services now enjoys a six times improvement in response time (from 12 seconds to less than two seconds), and SLAs are now consistently met. Another bonus: From the productivity view, representatives now handle 20% more customers in the same timeframe.

The move from VAX to AlphaServer systems was much easier than anticipated, the company claims. Between increased performance and server consolidation, the payback has been absolutely "tremendous." Moreover, as with most enterprises today, the Verizon Information Systems IT environment consists of diverse hardware and multiple operating systems. The interoperability of OpenVMS-based AlphaServer systems is important in this environment.

[16] This is not atypical. Many Compaq cluster environments run both the VAX and AlphaServer machines. This capability helps preserve investment and can be quite important in many environments.

CONSOLIDATION CASE STUDY THREE – SMALL BUSINESS SERVICE BUREAU

Founded in 1968, The Small Business Service Bureau, Inc. (SBSB) is a national small business organization with 50,000+ members. These members are self-employed, in partnerships, family businesses, home-based businesses, and owners involved in every kind of business or trade. Because of their small size, SBSB members share a need for products and services custom-designed for them at costs they can afford. Such products and services include health care coverage, dental, prescriptions, vision, hearing, disability, life insurance, retirement plans, business insurance, long term care, auto insurance, and more.

It is critical to the Worchester, Massachusetts-headquartered SBSB to keep its IT costs under control in the face of rapid growth. Such control contributes mightily to SBSB's ability to continue to offer best buys to its members. Long a VAX hardware and OpenVMS software-based firm, SBSB set clear goals for its AlphaServer upgrade and server consolidation effort. These goals included,

- Scalability to support an annual growth rate of 10 to 12%.
- Increased performance by separating SBSB's multiple databases from its application servers.
- Continuing to be able to leverage its long term investment and knowledge of OpenVMS with production-worthy, fully debugged, dedicated SBSB applications.
- Continuing to enjoy the ability of Compaq Global Services to provide maintenance and support for software and hardware.
- Save dollars at least in the six figure range from increased employee productivity, reduced floor space use, and increased performance.

According to SBSB Vice President of Information Technology, Richard Amico, these goals were met and then some. Organizations with existing VAX and OpenVMS environments that have similar goals can apply the SBSB story to their environment.

WHY SERVER CONSOLIDATION AT SBSB?

The outcome of SBSB's server consolidation effort is a two-node cluster consisting of a Compaq AlphaServer and a Compaq VAX, both running the OpenVMS operating system. SBSB started with multiple stand-alone VAXs running OpenVMS. The change enabled SBSB to cut its number of servers in half while delivering approximately three times the performance.

SBSB felt it was necessary to make the consolidation move not only for the reasons mentioned above but also because each day, as Amico puts it, "Literally, tens of thousands of transactions and call history data must be collected, managed, and retrieved instantaneously. In addition, our insurance quoting

engines are dynamic in nature and require updating and modifications on a regular basis." Clearly, a lot of horsepower was needed.

SBSB maintains a call center with sales and customer service representatives who handle thousands of daily outgoing and incoming calls. These 150 people support 50 to 70 different financial and insurance plans. The most modern technology was needed to keep all this running smoothly and cost-effectively.

All SBSB production systems (e.g., invoicing and accounts receivable) ran on Compaq platforms. The relationship between the two companies goes back to 1977. At that time, SBSB started with COBOL-based DECsystem computers. Happy with this infrastructure for 13 years, in 1990, SBSB moved and upgraded to a VAX platform running OpenVMS. Adding VAXes, the firm first used index files, then migrated to Compaq's RDb, and then to Oracle in 2000.

Needing more horsepower, SBSB added an AlphaServer in 1995. At that time one of the firm's VAX systems acted as the database server and application server. Soon overburdened, the response time of this gear was too slow. So the AlphaServer was brought in specifically to run the databases and the performance issue was removed. Today, 14 production databases that comprise some 50 to 55 GB are accommodated. Up to 120 GB may be handled by the existing Storageworks storage.

Before its recent server consolidation, SBSB ran a four-node cluster that handled both the applications and stored the databases. The consolidation to the two-node cluster running one VAX and one AlphaServer allowed SBSB to cut its number of systems in half while tripling performance. The cluster configuration today includes one VAX 7630 with three processors clustered to an AlphaServer 4100 system that holds all the databases. All of the production legacy applications run native on the VAX system, and both systems run OpenVMS v7.2. The savings from being able to continue to run the legacy production applications native on the VAX and the same operating system on both nodes are obvious.

Both the VAX to Alpha upgrade and the server consolidation brought major advantages to the Small Business Service Bureau. "Between the upgrade and server consolidation, we are seeing increased system performance, virtually unlimited scalability, greater employee productivity, and savings in floor space," says Amico. There are more benefits. "The new architecture easily supports an annual business growth in excess of 10%," Amico notes. "And it allows us to expand both the application and database servers independent of one another."

Employee productivity is still another benefit. "Our employees can process more applications because of the improved system performance," Amico observes. In fact, the current system is able to support almost three times the number of employees and three times as many transactions. The VAX to Alpha consolidation has also saved valuable floor space. Essentially, SBSB has increased its workload by 300% without having to increase the required floor space. System

management is another plus. "It certainly is easier to manage two systems rather than four or five," Amico says.

Dan Lebeau, Development Manager at the Small Business Service Bureau, adds another cost-of-ownership note of importance. "While software costs have gone up due to licensing more users, the system costs have gone down – so we have effectively been able to upgrade to these systems with no apparent increase in monthly lease costs."

Amico quotes real numbers for reduced costs. "When you combine such factors as faster performance, increased employee productivity, and reduced floor space, we estimate a total annual savings of approximately $200,000 – just by consolidating on AlphaServer systems."

SBSB planned its upgrades to take place over a weekend, so there was no service disruption. "The VAX to Alpha upgrade went quite smoothly," says Amico. "Compaq Global Services was instrumental in our upgrade process. We couldn't have done it without their help. They worked with us through the design, planning, and pre-implementation stages. Due to their thorough project planning and excellent attention to detail, the actual conversion occurred without a problem."

A LOOK BACK

The business requirements of large customer consolidation initiatives typically dictate a very secure, scalable, and ultra-high availability application platform infrastructure that has proven itself over many years of experience. The reason for this is because a single server going down may no longer affect only a small number of users or customers, but rather, many thousands in the new environment.

Compaq's customers are very satisfied with the outstanding clustering, partitioning, availability, and workload management capabilities of OpenVMS AlphaServers. Similarly, the OpenVMS AlphaServer family can readily meet other customer issues that offer substantial requirements for clustering, partitioning, availability, and workload management.

Developed as a result of many years of operating system and cluster experience, Compaq OpenVMS technologies such as, OpenVMS Galaxy software, OpenVMS Clusters, System Management, Continuous Computing, and other technologies are Compaq differentiators that work together to meet a wide variety of server consolidation needs.

When these differentiators are coupled with the capabilities of the state-of-the-art Compaq AlphaServer Series hardware, the server consolidation environments that can be handled in an industry-leading manner include homogeneous workload, heterogeneous workload, management/operational, application, and Oracle database consolidation. Described in detail above, these scenarios have been successfully implemented in the real world.