The HP OpenVMS Approach to High Availability Computing

Table of contents

invent

Table of contents (continued)


hp

The HP OpenVMS High Availability Mission Statement

HP will provide our customers with an environment in which their required applications and associated data are available **all the time** within specified performance guidelines regardless of hardware and software, and planned or unplanned occurrences.

The OpenVMS commitment to high availability involves continued support and enhancement of existing high availability and disaster tolerant features, the evolution of that environment based on growth requirements, and enhanced enterprise features such as OpenVMS partitioning, mixed architecture clusters[1], and hardware component hot swapping.

In addition, HP is committed to the extension of the OpenVMS operating system to the next generation of HP servers. OpenVMS has undergone a full port to the 64-bit Intel® Itanium®-based HP Integrity server line.

[1] AlphaServer systems with Integrity servers or VAX systems with AlphaServer systems.

1. Introduction

The business loss from service outages continues to grow rapidly. News clips of businesses losing millions of dollars of market capitalization in hours or days due to e-Commerce or operational outages are far too common. U.S. Bureau of Labor statistics and others show that for companies suffering traditional disasters, only 6% survive long-term, 43% never resume business, and 51% close down their business within two years.

Listed below are staggering statistics as reported from a 2001 survey conducted by Contingency Planning Research and Contingency Planning & Management, regarding hourly downtime. Of the IT professionals who responded:

- 46% stated each hour of downtime would cost their companies up to $50,000
- 28% said each hour would cost between $51,000 and $250,000
- 18% said each hour would cost between $251,000 and $1,000,000
- 8% said an hour of downtime would cost their companies more than $1,000,000

In the face of these negative financial impacts, IT professionals, working in concert with hardware and software vendors, battle back with high-availability solutions.

HP defines high availability as the ability of a computer system to provide access to applications and data — based on the level of services required by an organization and its users whenever and wherever its business needs dictate.

This white paper examines the capabilities that contribute to a high-availability environment and explores the industry-leading high-availability features of the HP OpenVMS operating system.

It should be noted that the OpenVMS operating system runs on three distinct server platforms: VAX systems, HP AlphaServer systems, and — incrementally rolling out at the end of 2004 and into 2005 — HP Integrity servers.

While high availability is a business-critical requirement in every industry, there are certain environments — healthcare, finance, manufacturing, telecommunications, and government, for example — where availability is the top priority. Because OpenVMS offers unparalleled availability, it has an outstanding presence in these industries through a broad range of high-availability solutions including products, services, and partnerships. At the same time, OpenVMS provides a core business-critical environment for in-house application development in any industry.

HP OpenVMS solutions combine industry-leading high-availability and disaster-tolerance capabilities with open database and application support. And the power of HP AlphaServer systems and HP Integrity servers delivers unrivaled flexibility and breadth of customer solutions.

2. What is High Availability?

High availability isn't something that can just be bolted on. The IT organization must approach high availability with a fundamentally holistic and systemic point of view. This means looking both at and beyond the IT environment including

- Physical infrastructure (HVAC, power, communications, isolation from extraneous disruptions, and so on)
- Policies and practices
- Hardware and software technologies
- Services (for technology and for general business continuity)

While the first two items — physical infrastructure, as well as policies and practices — are essential to a complete high availability environment, this paper will focus primarily on hardware and software technologies and, to a lesser extent on services. Considerable information on the first two can be found at www.hp.com/go/businesscontinuity.

Before discussing the technology itself, though, the question of how to measure availability must be answered.

Measuring Availability

While availability is popularly measured in terms of the percentage of a year's time that a system is up and running — the so-called "9s" — it is more precisely and more practically measured in terms of Recovery Time Objective (RTO) and Recovery Point Objective (RPO). The considerations and characteristics of each are presented below.

Recovery Time Objective (RTO)

- How soon after an event does the business process need to be available?
- Not all business processes need to be available at the same time
- Primary focus is transactions

RTO identifies how long your IT environment can be down before it affects the business. In other words, how much downtime is acceptable? Time in this case is measured forward from the point of the failure until normal processing can be resumed. This requirement may depend, too, on the cycle of the economic year. A retailer, for example, may be a lot more resilient to downtime in mid-August than in December's holiday consumer buying cycle. Contrast this with some environments, in which penalties are paid for every minute of downtime, around the clock.

Recovery Point Objective (RPO)

- How much work in progress can be lost?
- Not all work needs to be recovered to the same point
- Primary focus is data

RPO identifies the point at which the data is restored after a disaster, and thus quantifies the amount of data loss that is acceptable — if any. In this case, one measures backwards from the point of failure. Data vaulting, for example, may be all the protection an organization needs for certain parts of its IT environment while others may require recovery of even an in-flight transaction.

Once RTO and RPO are determined for a given environment, an organization can use this more granular measurement model to design and deploy finely grained high-availability hardware, software, and services solutions appropriate to specific locations, processes, time periods, and technologies.

Implementing High Availability

Hardware High Availability

One way to enhance availability at the hardware level is to cluster two or more computers together. In the most basic clustered environment, if one system goes down the other can assume the workload. In an OpenVMS clustered environment all systems function as though they are a single computer. Should one system fail, the workload is automatically distributed among the others with the result that users usually do not notice system failures nor is work interrupted or data lost.

High availability can be further enhanced through the use of multi-site clusters. In this scenario, each system or group of systems is in a geographically separate location. By splitting the computing environment among two or more locations that are at safe distances apart, an organization can ensure that it will remain in operation even in the face of having lost the use of a complete data center. By adding sufficient resources at the two sites, it may be possible to run most critical applications without significant degradation.

Software High Availability

Software plays a central role in creating a highly available environment. While the hardware can provide redundant paths and component failover capabilities, the ability of software to predict, detect, report, and react to failures is critical. To ensure availability, software must be able to transparently detect when a component has become unavailable, be able to quickly and concisely report the failure, and determine when another component can be used without losing system and application context.

Within a cluster, the software must be able to rely on other member systems to fulfill aspects of the application functions. If one of these resources becomes unavailable, the software must be able to find and use other available resources transparently and quickly. Because upgrades are a significant source of planned downtime, it is also essential that software upgrades are fast and easy.

Services are Key to High Availability

Services bridge the gap between what the technology provides and what the environment requires in terms of uptime. When the application environment (including hardware and software) is inherently highly available due to the capabilities of the technology, the need for emergency support services diminishes. In other words, when data is available and applications continue to run, support activities can occur at a convenient time, rather than in an emergency mode when the service or upgrade must be completed for the application to continue to function.

On the other hand, as an environment takes greater advantage of high-availability attributes of the technology, it becomes more important to develop and implement appropriate procedures, testing, and verification. Therefore, management of support services become vital; and change management becomes paramount.

Services can include consulting, which involves the design of a highly available environment; necessary simulations of current and future performance; and the detection of problems before they arise. The proper design of an environment — and the analysis of existing environments for robustness in the event of a change or failure — are all key aspects of ensuring a highly available system. Moreover, it is essential to document, monitor, and control all changes to the environment to avoid unintended consequences, errors, or degradation of RTO and RPO.

3. HP OpenVMS and High Availability

Design Theory

The OpenVMS operating system has been in existence for almost thirty years and continues to deliver the robustness necessary to run business applications reliably. While longevity is one measure of an operating environment's ability to be highly available, the inherent design of that environment is the true measurement. From its inception, OpenVMS was designed with availability, flexibility, and scalability in mind.

The design theory behind OpenVMS is that — regardless of the size or complexity of the hardware that is being used — applications should run in the same way on the OpenVMS operating system. In other words, an application that runs on a workstation could be moved to an enterprise-wide server and still run correctly. What's more, applications are able to access data either directly from the local system or remotely through connections with other systems. These connected systems can share resources for data access, system resources, and can be centrally or remotely managed.

OpenVMS runs on HP VAX, HP AlphaServer, and HP Integrity server systems. While each of these hardware platforms has inherent Reliability, Availability, and Serviceability attributes, OpenVMS and associated high-availability software products enable these environments to meet service-level requirements dictated by RTO and RPO. An overview of AlphaServer and Integrity server systems is available in Appendix 3.

OpenVMS System Management Tools

System management directly contributes to high availability. With HP OpenVMS, system management is easy because an entire cluster can be managed from a single workstation, either locally or remotely. This single-location management can significantly reduce management costs for large, complex systems while providing tools and performance data to improve cluster efficiency and reliability.

OpenVMS system-management tools reduce operator error and aid in keeping systems up and running efficiently. They help improve workload management among shared resources such as printers and disks, and provide a central interface through which third-party system management tools can operate. In addition, OpenVMS system-management tools maintain a history of system functioning and furnish easy and central manipulation of user accounts.

A complete set of system management solutions is available for OpenVMS including schedulers, console managers, event and performance advisors and analyzers, and others. These products are available from HP and from a number of third parties. For more information, visit www.hp.com/go/openvms.

OpenVMS Security

A major availability risk resides in the realm of security. Hackers, viruses, and worms, will adversely affect availability. Systems must have a level of security that supports the desired RPO and RTO. Table 1 below illustrates the HP focus on four key areas of security and the OpenVMS features that contribute to high availability.

Table 1: OpenVMS = Secure by Design

Integrated Security Model since V1.0

- Always C2 Compliant
- Virus resistant architecture

Network Security	Data Integrity and Confidentiality	Intrusion Detection and Analysis	Access
Supports industry-standard network security tools: • Kerberos • OpenSSL • SSH • IPSEC	Supports industry-standard encryption: • CDSA • GnuPG • Encryption for OpenVMS (securing backups)	• Host-based cluster-wide intrusion detection • Comprehensive audit analysis and reporting	• Single security domain across host or cluster • Support for multiple authentication methods • Fine-grained privileges • OpenView integration

What's more, data assembled from primary sources by TechWise Research indicates that the security risks with OpenVMS are the lowest in the industry as shown in Table 2.

Table 2: Comparison of security risk

	Negative Effects		
Incidents	**OpenVMS**	**Solaris**	**AIX**
Average Number of cluster crashes per year caused by a virus or worm	0.28	2.84	1.47
Average annual hours of cluster downtime due to a virus or worm	0.88	4.32	6.73
Number of security advisories issued by CERT© Coordination Center between 2000 and 2003	2	29	28
Common Vulnerabilities and Exposures (CVEs) identified by MITRE corporation	6	167	89

Unlimited OpenVMS Growth

HP offers virtually unlimited OpenVMS growth for years to come through a clear, cost-effective upgrade path with hardware, software, services, and option upgrades. Upgrades to SMP and high-performance clustering provide effective scaling for major performance increases, while flexibility is further assured by extensive compatibility with industry standards.

Beginning with OpenVMS Version 8.2 there is a single-source-code stream for both the HP AlphaServer system and HP Integrity server platforms. This ensures that all non-hardware-dependent enhancements will be released on both server environments in the same timeframe including high-availability and disaster-tolerance capabilities.

4. HP OpenVMS High Availability Features

The previous section provided an overview of the basic qualities that make OpenVMS a natural choice for high availability. This section will focus on specific OpenVMS capabilities that contribute to building high-availability and disaster-tolerant environments.

The OpenVMS File System

The OpenVMS File System provides for I/O operations performed in carefully designed sequences, and is intended to provide for consistent and recoverable data even if a related sequence of disk I/O operations does not complete.

These careful updates greatly reduce exposure to file system corruption, both on a stand-alone system and in an OpenVMS Cluster. Integrated support for two-phase commit transaction management as well as journaling capabilities further ensure that applications have clear and consistent data available at all times. This combination of file system features preserves critical application data over system and even application failures.

In conjunction with products such as Reliable Transaction Router (RTR), the two-phase commit transaction processing capabilities of OpenVMS can be extended across separate clusters, creating a software-based, global, transaction fault-tolerant environment[2].

OpenVMS configurations provide for the shadowing of disk storage, permitting application data to be transparently mirrored across available storage hardware, regardless of its location. Volume Shadowing, or host-based shadowing, fully interoperates with controller-based RAID. This permits transparent data redundancy — with local or remote storage. Both the file system and Volume Shadowing provide transparent recovery from common errors, reducing the effort and the risks involved with manual operator intervention.

[2] For more information on Reliable Transaction Router refer to www.hp.com/go/rtr.

Integrated and Distributed Batch Processing

HP OpenVMS provides integrated and distributed batch processing. Batch processing permits non-time-critical applications to be scheduled in the background and processed on any of specified sets of available systems. OpenVMS also provides for batch restart — permitting batch jobs to checkpoint application data and automatically restart after a system shutdown, failure or reprioritization. This provides a simple way to schedule non-priority tasks to gather available resources across a collection of nodes, or to schedule high-priority tasks transparently and automatically, without regard for which specific nodes are available when the job runs.

Recovery from Failure

As important to availability as reliability and uptime, recovery from failure helps maintain availability. HP OpenVMS provides integrated and optional capabilities — for such tasks as operator communications, DECevent error handling and analysis, and system environment monitoring — permitting systems to be repaired and brought back online quickly. With failure prediction, problems can be detected and reacted to before they become critical. Losing one of a redundant pair of power supplies isn't a problem — until the other is lost.

Similarly, the rapid detection and response to security attacks can reduce exposure, and can be used to identify and repair any breaches that might occur. This gets systems back online quickly and helps keep them online. OpenVMS provides for security alarms as well as for the auditing of system activity.

And as a last resort, OpenVMS provides an integrated backup utility and associated calling interface.

OpenVMS Partitioning

Those who use OpenVMS partitioning in an AlphaServer environment will be familiar with the term "Galaxy" by which it is known. Similar functionality is also being developed for OpenVMS on HP Integrity servers.

OpenVMS partitioning enables the running of multiple instances of OpenVMS in the same system. These instances can be reconfigured dynamically. And these instances are peers — operating in parallel and with direct access to the hardware. Each OpenVMS partition instance boots and shuts down independently.

An OpenVMS partitioned computing environment provides exceptional scaling in symmetric multiprocessor configurations. Further, OpenVMS partitioning capabilities provide flexibility around the assignment of system resources across the instances — to balance continuously changing system loads; meet application testing or upgrade requirements; or meet organizational configuration requirements.

Instances in an OpenVMS partitioning computing environment can be clustered with other instances within a partitioning system using the shared-memory clustered interconnect to communicate with each other. Instances in one OpenVMS partitioning system can also be clustered with instances in another OpenVMS partitioning system or with cluster nodes in non-partitioning systems. Instances clustered outside of a partitioning system use traditional cluster interconnects.

With AlphaServer systems, portions of memory can be shared among processes and among processes running in multiple instances within an OpenVMS partitioning system. This permits applications to share data and to share data caches, and even to shut down and restart the operating system without having to reload the application data caches. Even if OpenVMS partitioning instances are not configured as members of the same AlphaServer OpenVMS Cluster, this backplane-speed shared memory is available.

This flexibility permits an organization to choose the OpenVMS partitioning configuration that best meets its reliability and availability needs. And a configuration can be changed and even upgraded dynamically — all while critical applications remain running.

OpenVMS Clusters

HP OpenVMS Cluster software is an integral part of the OpenVMS operating system, providing the basis for many of the key capabilities used by OpenVMS enterprise solutions. A full "shared everything" cluster design that has been in existence for more than 20 years, OpenVMS Cluster software allows for the maximum in expandability, scalability, and availability for mission-critical applications. In fact, a DH Brown Associates study[3] of IT managers with OpenVMS cluster experience published in June 2004 reported that

". . . the consensus was that OpenVMS clusters remain "superior" in many senses to all of the UNIX-based clusters available today in terms of sophistication of implementation (e.g., clustering integrated into file system, users can connect to any server in the cluster, clustering built into the operating system, up to 96 nodes in datacenters up to 800 km apart). Moreover, they were aware (to varying degrees) that an OpenVMS cluster can simultaneously be a scale-out, high availability, and disaster-tolerant cluster depending on the scripts used, the interconnect inter-nodal distances, and the application or applications deployed."

[3] Appendix 2 contains a more extensive excerpt from the study.

A key enabler of the OpenVMS role in the HP Adaptive Enterprise, OpenVMS Cluster software has been an industry bellwether technology[4] where enterprise environments must adapt to growth and changes over time. These capabilities also make the total cost of ownership (TCO)[5] for an OpenVMS application environment extremely attractive.

A Closer Look at OpenVMS Cluster Software

An HP OpenVMS Cluster is a highly integrated organization of VAX and HP AlphaServer systems — or AlphaServer and HP Integrity server systems-applications, operating systems, and storage devices. These systems can be connected to each other and storage components in a variety of ways, depending on an organization's needs.

Figure 1 shows a typical small OpenVMS clustered environment. The environment is characterized by

- The ability to have multiple server architectures running OpenVMS
- Common/shared storage environments — Fibre Channel and/or CI
- Non-shared storage but the ability to serve the data to any node in the OpenVMS cluster
- The ability to share resources freely within the cluster

[4] "…the Disaster Tolerant OpenVMS Cluster. Still considered the "gold standard" in commercial clustering, it's not unusual for OpenVMS Cluster uptimes to be measured in years." (David Freund, Illuminata, 2002)

[5] See http://h71000.www7.hp.com/openvms/whitepapers/ for independent analysis results by TechWise Research Inc.

Fig 1. Basic OpenVMS Cluster configurations

Note: For the purposes of development and migration, customers can have VAX systems in their mixed-architecture Integrity clusters. However, should a problem be discovered because of the inclusion of the VAX systems, the customer will need to remove those VAX systems. In a production environment, VAX systems are not formally supported in a mixed-architecture Integrity environment…

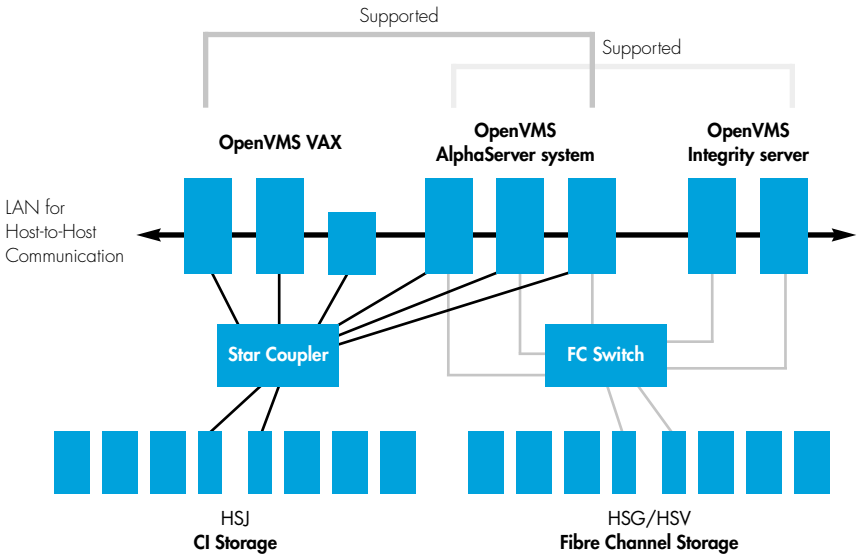

OpenVMS Cluster systems provide the ultimate in a highly available, scalable, and flexible computing environment. The cluster also allows the connection of systems of all sizes and capacities to achieve an easy-to manage, single virtual system. And the inherent capabilities of OpenVMS clustering allow systems to be connected at distances from just a few inches up to 800 KM apart.

As OpenVMS moves to the Integrity server line, the role of clustering to allow customers a seamless integration into their environment cannot be understated. Due to its inherent mixed-architecture clustering capability, OpenVMS clustering allows both AlphaServer systems and Integrity servers to be mixed within the same cluster environment.[6]

Fibre Channel SAN storage can be shared between the two architectures freely, given that the on-disk file structure is identical between the two. Additionally, applications can be run on either architecture within the cluster and allow users throughout the cluster to use the application as necessary.

[6] Incorporating HP Inegrity Servers into an OpenVMS AlphaServer system environment can be as straightforward as adding one or more nodes to the cluster.

Disaster Tolerance

Disaster tolerance is the ability to withstand the functional loss of an entire data center arising from events such as natural disaster, sabotage, loss of power, or human error. IT environments are typically protected against such catastrophes by extending cluster and data-replication capabilities beyond the local geographic area. This might be across a campus, the other side of the world, or anywhere in between.

The OpenVMS high-availability capabilities allow for this. Indeed, OpenVMS clusters can be separated by as much as 800 kilometers between nodes while still functioning as a fully operational single-system environment. And this occurs even if one part of the environment becomes completely incapacitated. Moreover, data replication has been qualified up to 90,000 kilometers.

Considerations that go into designing and implementing a multi-site cluster include:

- Identifying critical data for shadowing (replication)
- Site selection
- Communication among the sites
- Policies and procedures regarding failure and restart of either site
- Appropriate setting of system parameters to enable smooth failovers

Moving from High Availability to Disaster Tolerance

Section 2 of this paper discussed the measurement of availability using RTO and RPO. It pointed out that this measurement model delivers a more granular picture of a given environment, enabling an organization to more appropriately apply resources to specific locations, processes, time periods, and technologies. Implementing OpenVMS capabilities along a continuum from high availability to disaster tolerance can be equally as granular.

This means that once an organization has precisely defined its RTO and RPO requirements, it can, with equal precision, build an OpenVMS environment to meet those requirements. And — because of OpenVMS scalability — it can more accurately purchase the level of computing resources it needs at present without boxing out future growth. In the end, this contributes to a lower TCO, ease of use, and more facile change management. The next few pages present an overview of the incremental steps that an organization can take to move easily and economically along the OpenVMS high-availability/disaster-tolerant continuum.

Figure 1 on the previous page, showed a basic local high-availability cluster environment. Here, Figure 2 depicts an OpenVMS cluster design with two local high-availability clusters connected over a distance to provide disaster tolerance.

This is not only a disaster tolerant configuration but is, in essence, a four-node cluster with two pairs of two nodes in two different locations. Host-based volume shadowing provides data protection.

Fig 2. Two local high-availability clusters connected in a DT configuration

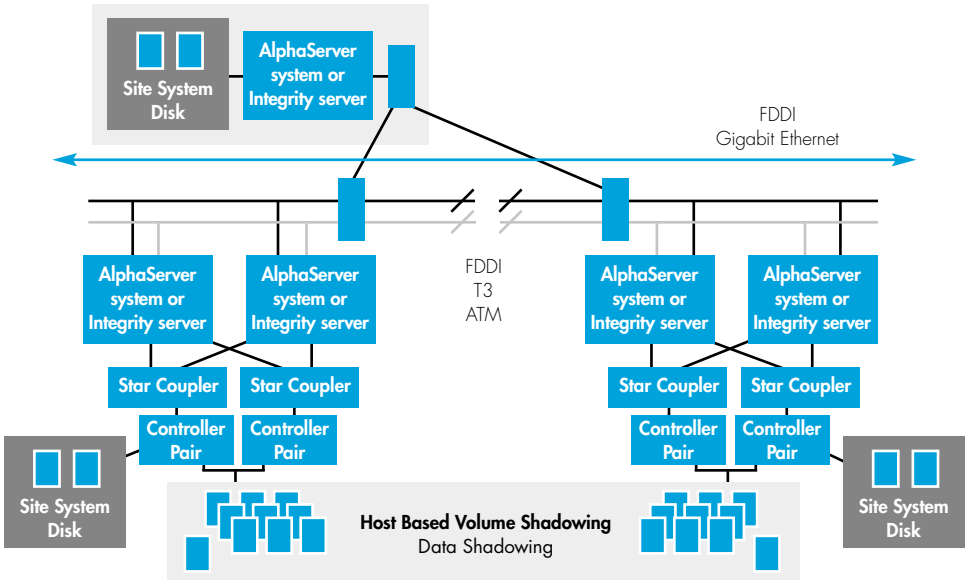

Fig 3. Quorum disk

- keep quorum equal on each site
  - uneven quorum risks "creeping doom"
- use quorum site or out-of-band quorum adjustment
  - availability manager or DTCS management tools
- use boot time commands to implement single-site booting

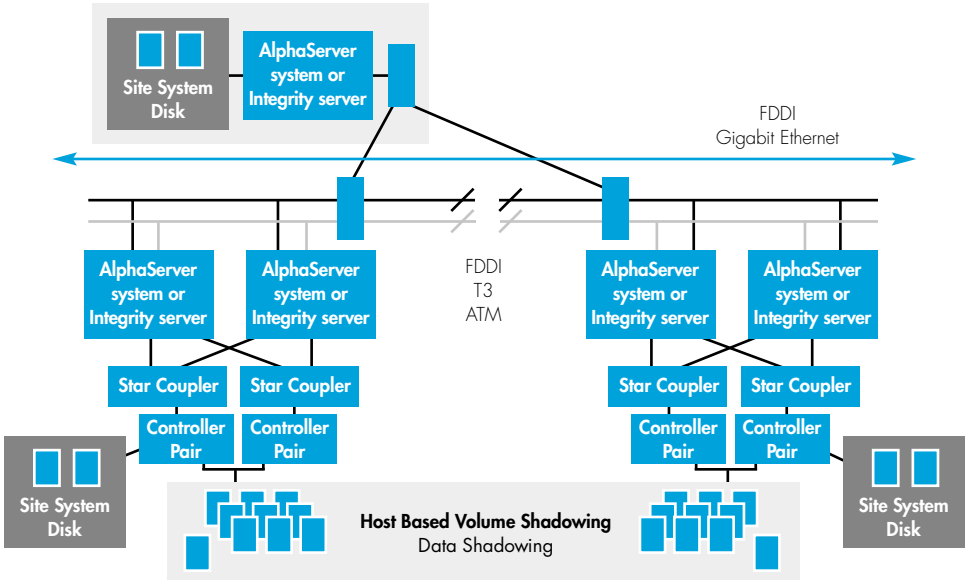

Figure 3 introduces the quorum disk that provides mediation between cluster nodes. Without this mediation a condition known as "split-brain" syndrome can arise. If a link is broken between sites, each site operates independently. When reconnected, each site will have a unique version of data and transactions and the system will not know how to combine them. A quorum disk (or quorum node) will mediate and determine the proper actions to take on both failure and resynchronization.

Fig 4. Fibre Channel direct data storage

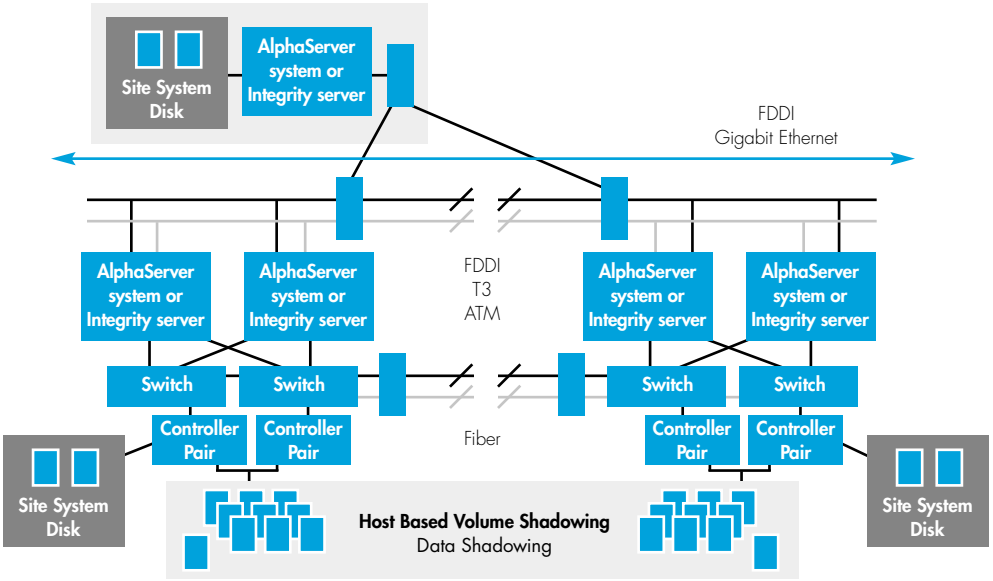

Figure 4 is similar to the configuration shown in Figure 3 but it adds Fibre Channel direct data storage to the data storage link for improved data replication performance.

Fig 5. Host-based volume shadowing

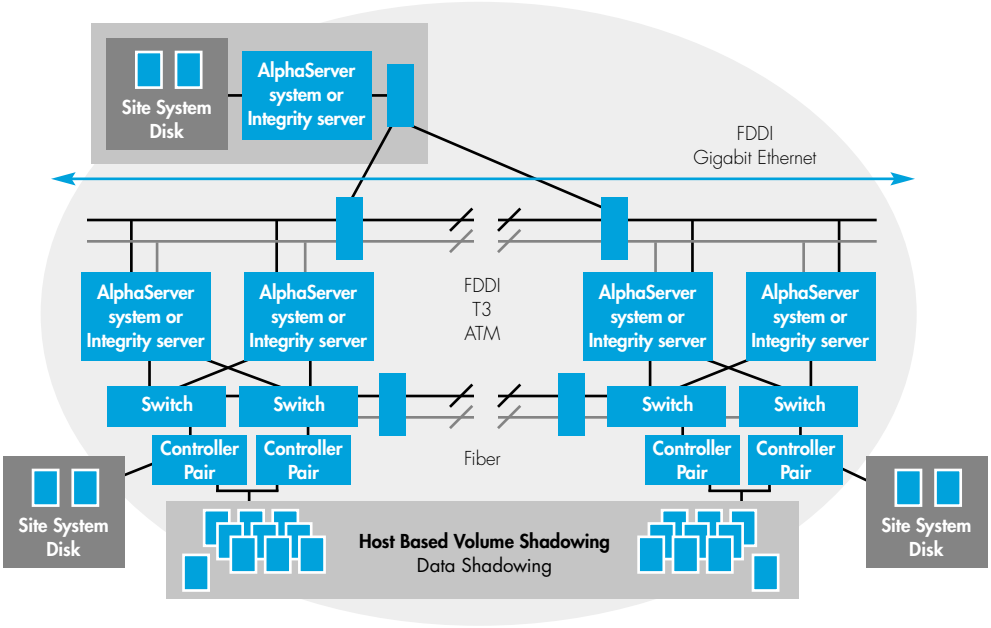

Figure 5 builds on the above configuration with the capability to copy system disks with host-based volume shadowing.

Fig 6. Data Replication Manager

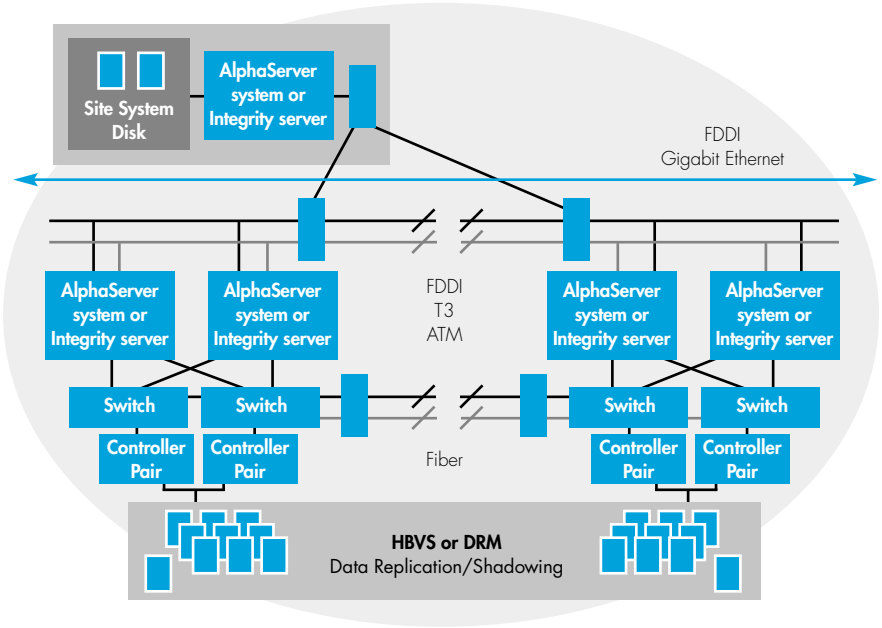

The design in Figure 6 includes continuous-access Data Replication Manager (DRM) to copy system disks. DRM is an ideal solution for mirroring data online and in realtime to remote locations via a local or an extended Storage Area Network (SAN). Using DRM software, data replication is performed at the storage-system level and in the background to any host activity.

Fig 7. Replication over distance

This chart shows the throughput at different distances using 1 lun for the testing – then we stopped testing!

Note that the "cold standby" systems can also be doing other work.

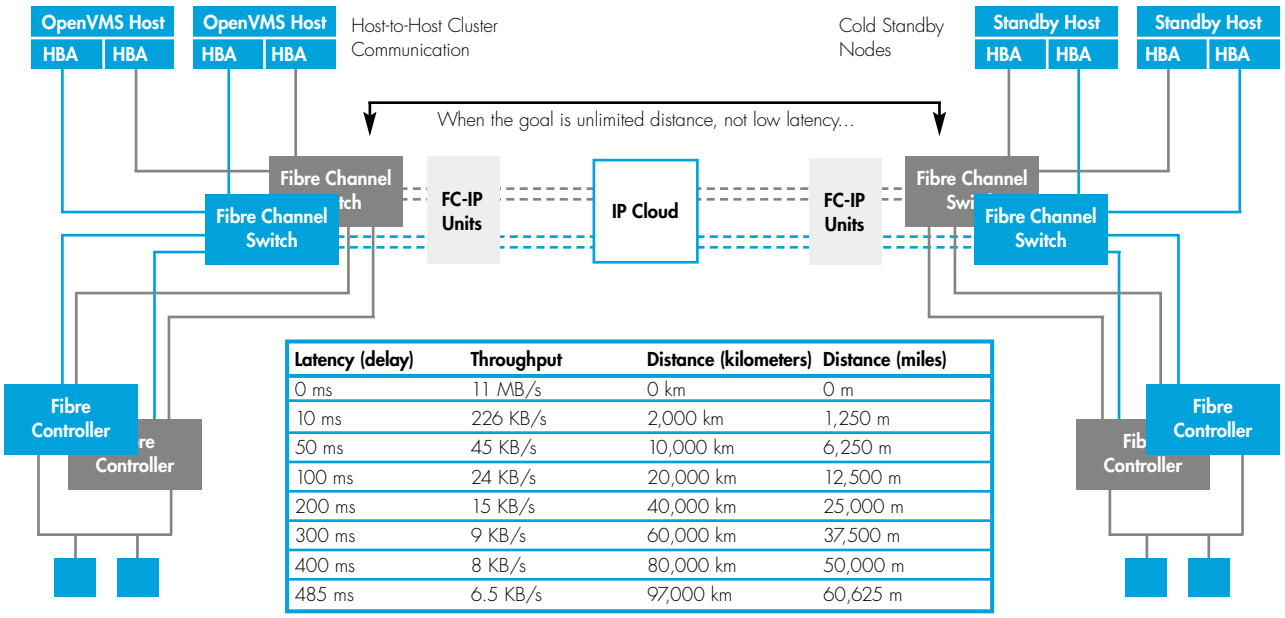

Latency (delay)	Throughput	Distance (kilometers)	Distance (miles)
0 ms	11 MB/s	0 km	0 m
10 ms	226 KB/s	2,000 km	1,250 m
50 ms	45 KB/s	10,000 km	6,250 m
100 ms	24 KB/s	20,000 km	12,500 m
200 ms	15 KB/s	40,000 km	25,000 m
300 ms	9 KB/s	60,000 km	37,500 m
400 ms	8 KB/s	80,000 km	50,000 m
485 ms	6.5 KB/s	97,000 km	60,625 m

With disaster-tolerant configurations, the question arises of how far away from the primary site can data be stored? The real question that needs to be answered is "How "realtime" must the data be at the remote site?"

As shown in Table 3, the results of tests performed by HP DRM Interconnect Engineering, indicate that OpenVMS, right out of the box, can achieve data replication thousands of miles away. The circumference of the world is only 26,120 miles (40,820 km) and testing was stopped at simulated distances more than twice that (60,625 miles or 97,000 km). So, achieving data replication at great distances isn't really the problem. However, whether this approach works depends on the problem to be solved.

Table 3. Data Replication Distances

Latency (milliseconds)	Throughput (seconds)	Distance (kilometers)	Distance (miles)
0 ms	11 MB/s	0 km	0 m
10 ms	226 KB/s	2,000 km	1,250 m
50 ms	45 KB/s	10,000 km	6,250 m
100 ms	24 KB/s	20,000 km	12,500 m
200 ms	15 KB/s	40,000 km	25,000 m
300 ms	9 KB/s	60,000 km	37,500 m
400 ms	8 KB/s	80,000 km	50,000 m
485 ms	6.5 KB/s	97,000 km	60,625 m

The left side of Figure 7 shows an example of an OpenVMS Cluster using Fibre Channel switches as a storage interconnect. It also illustrates the deployment of Fibre Channel over a network to replicate data for disaster recovery. The systems on the top right are "cold standby nodes" and are not active members of the cluster on the top left. They are there for recovery. The efficiency of this environment would depend on the management of the write operations. For example, using journal files on a local disk and determining when and how much data must be updated through "mini-copy" operations can impact how much latency your operation can withstand and still provide the desired level of performance.

Economists often talk about a concept known as "Opportunity Cost". The true cost of something, economists say, is what you give up to get it. Everything you do has an opportunity cost. And, when it comes to electronic data interchange across great distances, another important question needs to be answered: "What problem are you trying to solve?"

Figure 7 illustrates that a data recovery strategy could be deployed over extreme distances. At system start-up, of course, processing time would have to be invested to create a duplicate copy of the entire local database at the remote site. Once fully operational, "mini-copies" could be used to keep the two databases consistent.

With device journaling for database recovery, each journal file contains a transaction applied to the database. The journal files can be used to recover a database several ways. One way would be to reconnect the distant members, and then replay all of the transactions, using the journal files. Another way would be to restore the database from backup and apply the journal files to the database backup. This would require the restoration of the remote devices as well.

5. The Continuous Enhancement of HP OpenVMS

HP continues to enhance an already robust set of high-availability and disaster-tolerant features with offerings contained in each new OpenVMS release.

OpenVMS V7.3-1

OpenVMS V7.3-1 delivered a set of enhancements centered on improving Fibre Channel storage performance and availability. Areas of improvement included performance enhancements within the operating system, as well as improved failover capabilities with support of Multipath Tape capabilities. Additionally, the OpenVMS clustering capability was enhanced to allow the serving of data from one cluster node to another using the Fibre Channel storage paths instead of the traditional host-to-host links. The combination of these features delivered greater availability and much improved DT performance.

OpenVMS V7.3-2

OpenVMS V7.3-2 delivered a set of key improvements in the area of Host-Based Volume shadowing, and cluster-wide device access. Starting with V7.3-2, customers are now able to dynamically increase the size of a storage volume, up to a maximum size of 1 TB of logical space. This feature enables them to take advantage of EVA-based storage solutions, specifically the virtualized capabilities therein.

In addition, Host Based Volume Shadowing (HBVS) was improved to allow members of a shadow set be of a different size. This allowed customers to build up to a three-member mirror set with any device that was available, with the smallest member controlling the overall shadow set size. However, when this feature is combined with the Dynamic Volume Expansion feature, the size of a shadow set can be dynamically increased by replacing set members with different devices, and dynamically increasing the member size one by one until the shadow set has increased in size - without taking the mirrored set down!

Late in 2004, a feature was made available for OpenVMS V7.3-2 to allow a minimum MERGE (or Minimerge) to take place for Fibre Channel storage within a shadow set. With CI-based storage (HSJ/HSC), the storage controller kept a physical log of write changes to a device if it was removed from the shadow set. When the member was introduced into the cluster again, only a MERGE of the changes was required. With Fibre Channel, no such controller-based logging was available. OpenVMS then introduced this new Minimerge feature for V7.3-2, which allowed any FC-based storage to have a bitmap log kept within the OpenVMS host to track write changes, such that when a shadow set member was brought back into the cluster, only the write changes were MERGED. This feature is enabled via a remedial kit for V7.3-2, and is included with the V8.2 code as released.

OpenVMS V8.2

OpenVMS V8.2 is the first release of OpenVMS to support the industry-standard HP Integrity server family based on the Intel® Itanium® 2 architecture. This release contained a limited set of new features for Availability and DT (the Minimerge feature listed above). However, the main capability delivered with this release was the inherent ability to have both AlphaServer and Integrity server systems operate in the same OpenVMS Cluster.

These two platforms coexist in the same cluster, share the same FC SAN infrastructure, and can connect using the same LAN-based interconnects used with OpenVMS clusters in the past. This allows an OpenVMS user to have a mix of AlphaServer systems and Integrity-based systems as part of the same cluster, separated to the limits of OpenVMS clustering, and use the same data at any location. Features such as HBVS, Clusterwide File System, and RMS work freely across both environments, and failover from one location to another continues to be seamless.

6. Conclusion

Industry Solutions with Business Partners

HP will continue to strengthen support of OpenVMS horizontal solutions through business-critical solutions partnerships with, for example, Oracle and major system management, development, and business application companies. At the same time, HP is also increasing investments to provide tightly integrated Polycenter replacement products.

Working together with its partners, HP will continue to enhance vertical market segmentation solutions in which we provide leadership in integration, management solutions, and more. In addition to providing reliable, dependable core OpenVMS business-critical computing environments that meet the needs of any industry, HP will focus on the following vertical markets:

- Financial Services
- Manufacturing
- Telecommunications
- Healthcare
- Government

The bottom line is that HP will continue to provide customers with an environment in which their required applications and associated data are available **all the time** within specified performance guidelines regardless of hardware and software, and planned or unplanned occurrences.

Appendix 1

Additional resources and information on OpenVMS high availability

For more information on the high availability solutions offered by HP OpenVMS, please see our Web site at: www.hp.com/go/openvms*.

To see an interactive demonstration of OpenVMS high availability and disaster tolerance go to http://h71000.www7.hp.com/availability/disaster.exe

Appendix 2

Appendix Table 1. A table summarizing OpenVMS high-availability features, functions, and benefits

Feature	Function	Benefit
1. 64-Bit Addressing	Provides the user application up to 8 TB of address space.	Allows huge amounts of data to be accessed in memory.
2. Batch Checkpoint and Restart	Support for batch procedures to checkpoint activity for potential later restart.	Enables the restart of batch jobs from the last checkpoint.
3. OpenVMS Clusters	Let 96 or more nodes share storage resources, files, individual records, and a wide variety of resources. Distances of hundreds or thousands of kilometers can separate member systems.	Allow physically separate systems to share resources and cooperate.
4. Device Flexibility	Support for a large number of storage and communications interconnects permits configuration flexibility.	Permits great flexibility in the configuration and allows the tailoring of configurations to requirements and budget.
5. Distributed Lock Management	Provides distributed coordination, control, arbitration, and cluster communications. Permits elections of primary or unique applications, detection of application failure, and the rapid communication of data items such as sequence numbers.	Simplifies programming, improves reliability, and greatly eases coordination tasks.

Appendix Table 1. A table summarizing OpenVMS high-availability features, functions, and benefits (continued)

Feature	Function	Benefit
6. Environmental Monitoring	Packages and interfaces available to monitor power supply status, system thermal, and various other characteristics.	Permits automated local or remote monitoring and recovery.
7. Partitioning	Provides for multiple instances of OpenVMS in the same system. Configurations ranging from stand-alone to fully clustered are permitted. Permits application and system scaling in large multiprocessor configurations.	Permits different workloads to operate in parallel, each optimized for its own performance and configuration flexibility.
8. Global Sections	Portions of the address space can be shared among processes on the system, and among processes across OpenVMS Galaxy instances.	Simplify programming, resulting in better application reliability and performance.
9. Host-Based Shadowing	Provides for fully distributed device mirroring across local storage and across hundreds of kilometers.	Avoids single-site failures.
10. Integrated BACKUP Utility	Permits the creation and the reliable restoration of system, application, and file backups, and provides an API.	Eases instituting and integrating recovery procedures. Reduces MTTR.
11. Integrated Batch Support	Choose to run interactive tasks directly when required, and to schedule batch activities for later processing when system resources are available, on any node in an OpenVMS Cluster.	Allows simpler management; configuration flexibility; the ability to schedule background and maintenance tasks; and schedule tasks to transparently soak up spare processing when available.
12. Integrated Debugging	Applications can dynamically activate the debugger (even debugging themselves), and issue commands to help locate and resolve application problems, speeding problem resolution. Manual and remote operations are also fully supported.	Speeds problem resolution through easier development and debug and support, and better application error reporting.
13. Integrated Error Logging	Integrated DECevent error logging and reporting helps identify failing components immediately.	Reduces MTTR.
14. Integrated Event Notification	Permits applications to detect new or departing members of an OpenVMS Cluster, and to track many other system events.	Simplify design and programming, integrates management.
15. Integrated Networking	DECnet, IP, ICC, and various other network protocols available.	Multi-site, multiplatform, multipath. Avoids platform- and site-specific failures.
16. Integrated Operator Communications	Permit system users and applications to communicate with system operators.	Provide simpler and integrated management.
17. Integrated Transaction Processing	Full two-phase commit support. Permits applications to perform all necessary tasks in a transaction, or none of them. Works in conjunction with RMS Journaling and various other components and packages.	Provides a reliable single and common interface for transaction environments. Ties in with packages such as Reliable Transaction Router (RTR).
18. Integrated Security and Auditing	Prevent and track security-relevant activities. An OpenVMS Cluster is a single security domain.	Enables the easy location, prevention, and resolution of application, internal, and external security problems.
19. Minimerge	Enables the merging of only the physical disk changes, rather than a full merge.	Minimerge is significantly faster than a full merge operation. It merges only those areas of the shadow set where write activity is known to have occurred.

Appendix Table 1. A table summarizing OpenVMS high-availability features, functions, and benefits (continued)

Feature	Function	Benefit
20. Remote Management Capabilities	The OpenVMS Management Station permits remote management, operation, and control of OpenVMS systems.	Simplify management of multiple (potentially disaster-tolerant) sites.
21. RMS Journaling	Mechanism for recovery from file-level, system, or application failures.	Permits recovery from application failures, various programming errors, and system failures.
22. RMS and XQP	Provide common file formats, common tools, and record-level file sharing across multiple applications, across multiple cluster members, and remotely using DECnet.	Highly reliable and proven file systems. Ease debug and recovery and provide common tools available to all applications.
23. Rolling Upgrades	OpenVMS Cluster configurations can run multiple versions of OpenVMS, and can continue to share resources.	Enables the phase-in of fixes and new application releases, and perform OpenVMS upgrades without shutting down the OpenVMS Cluster. Better flexibility reduces risk, reduces downtime.
24. Shared Storage	OpenVMS Cluster members share read and write access to storage devices, permitting flexible and distributed configurations. This includes the ability to share system disks among systems.	Reduces management overhead, simplifies configuration, reduces application design, development, and maintenance costs. Enables rolling upgrades.
25. Volume Rebuilds	Automatic and transparent rebuilding of volume file structures and shadow-set members.	Reduce management. Provide early and automatic detection of problems.
26. Web Servers	Netscape FastTrack and other Web servers available. Java and Perl are also available.	HTML files and CGI scripts can be shared across cluster members as well as across the Internet.

Appendix 3

An overview of HP 64-bit server platforms

HP AlphaServer systems and HP Integrity server systems combine high-performance processors with innovative technologies including high-performance interconnects and symmetric multiprocessing (SMP).

They offer full 64-bit addressing and fast 64-bit I/O with Very Large Memory (VLM) and Very Large Data Base (VLDB) capabilities. HP has years of experience bringing the power of 64-bit computing to high-availability solutions. Version 8.2 of OpenVMS, released in December 2004, runs on both AlphaServer and Integrity server systems. When combined with OpenVMS Clusters, AlphaServer systems and HP Integrity server systems combine performance with the high-availability capabilities needed to run business-critical applications.

Integrating HP Integrity servers in an OpenVMS Environment

Planning and implementing the integration of HP Integrity servers into an AlphaServer OpenVMS environment is similar to what one would do to plan an expansion or upgrade with AlphaServer systems. Indeed, adding an Integrity server system to an OpenVMS environment can be as straightforward as introducing a new node to the cluster.

In stepping through the planning and implementation process it is important to keep in mind that, all things being equal, an organization can opt for either AlphaServer systems or Integrity servers for quite some time to come. HP support for the AlphaServer system family extends well into the future. From a business perspective, though, it is prudent to explore the possibilities that exist with the industry-standard Itanium® architecture with an incremental integration of Integrity servers.

A Wide Range of Performance Choices

With very fast processors, lots of memory, and all the 64-bit components to take advantage of them, HP systems offer an enormous boost in database and application performance and scalability. HP is the only company that can bring this advantage to customers.

Application performance increases and application features become more robust throughout the HP AlphaServer system and HP Integrity server families. AlphaServer systems range from the one- or two-processor DS series through the ES series, and on up through the GS series

which offers up to 64 of the latest EV7z processors and 512GB of memory. Similarly, the HP Integrity server family ranges from the one- or two-processor rx1620-2 up through the multi-cell, 128-processor, 1-TB-memory Superdome.[7]

OpenVMS and HP server technologies such as clustering, software and hardware RAID, and distributed resource coordination all allow the hardware and the application to provide the most robust system possible.

Appendix 4

IT manager opinions of OpenVMS Cluster technology

The following is quoted with permission from *The United States User View of Business Continuity and Disaster Recovery*, DH Brown Associates, June 2004

OpenVMS Clusters Set the Pace

The seasoned IT managers with OpenVMS cluster experience highly appreciated this technology's unique features. They remembered that OpenVMS clusters originated the cluster generation with dedicated hardware and software. They were also cognizant of the fact that these systems have evolved to serve a wide variety of applications and environments.

In fact, the consensus was that OpenVMS clusters remain "superior" in many senses to all of the UNIX-based clusters available today in terms of sophistication of implementation (e.g., clustering integrated into file system, users can connect to any server in the cluster, clustering built into the operating system, up to 96 nodes in datacenters up to 800 km apart). Moreover, they were aware (to varying degrees) that an OpenVMS cluster can simultaneously be a scale-out, high availability, and disaster-tolerant cluster depending on the scripts used, the interconnect inter-nodal distances, and the application or applications deployed. These applications range from custom-written customer applications (most of the individuals interviewed had these) to general-purpose uses of the Oracle database. One interviewee noted that the ability to port an OpenVMS application from a single node to a cluster with minimal effort had served him well. The consensus of the interviewees was that OpenVMS clusters served their purposes well. This was deemed to be especially true when a general-purpose UNIX cluster would not do at all. For these interviewees, it was all a question of what capability was needed and what the cost of that capability would be. They did note that they thought that in the long run, there would be less of a need for OpenVMS cluster capability at the high end as UNIX clusters continue to improve and encroach upon this space. HP has a superb opportunity here with the forthcoming Itanium-based OpenVMS clusters (e.g., on Superdome). This solution features a superior cluster architecture running on a high-end chip. It also provides optimal performance to those high-end applications that require what the OpenVMS cluster offer. The challenge here is to get the message out in the face of all the other cluster offerings on the market, including those from HP. An additional challenge is to overcome the perception that OpenVMS is an "old" technology despite its superior technical features.

[7] As of publication OpenVMS is qualified on the rx1620-2, rx2620-2, and rx4640-8 Integrity Servers. By the first half of 2006 it will be qualified on the entire entry-level, mid-range, and high-end HP Integrity server family. OpenVMS will also be qualified on HP Integrity blade servers.

To learn more about HP's offering, visit **www.hp.com**.

5983-0191EN, 12/2004