


*EV7 AlphaServers Deliver  
Enhanced RAS and  
Powerful Performance*

*June 2003*



*A D.H. Brown Associates, Inc. White Paper Prepared for  
Hewlett-Packard*

This document is copyrighted © by D.H. Brown Associates, Inc. (DHBA) and is protected by U.S. and international copyright laws and conventions. This document may not be copied, reproduced, stored in a retrieval system, transmitted in any form, posted on a public or private website or bulletin board, or sublicensed to a third party without the written consent of DHBA. No copyright may be obscured or removed from the paper. D.H. Brown Associates, Inc. and DHBA are trademarks of D.H. Brown Associates, Inc. All trademarks and registered marks of products and companies referred to in this paper are protected.


This document was developed on the basis of information and sources believed to be reliable. This document is to be used "as is." DHBA makes no guarantees or representations regarding, and shall have no liability for the accuracy of, data, subject matter, quality, or timeliness of the content. The data contained in this document are subject to change. DHBA accepts no responsibility to inform the reader of changes in the data. In addition, DHBA may change its view of the products, services, and companies described in this document.

DHBA accepts no responsibility for decisions made on the basis of information contained herein, nor from the reader's attempts to duplicate performance results or other outcomes. Nor can the paper be used to predict future values or performance levels. This document may not be used to create an endorsement for products and services discussed in the paper or for other products and services offered by the vendors discussed.

# TABLE OF CONTENTS

<b>EXECUTIVE SUMMARY .....</b>	<b>1</b>
<b>INTRODUCTION.....</b>	<b>2</b>
<b>KEY RAS IMPROVEMENTS.....</b>	<b>4</b>
ON-CHIP L2 .....	4
RAID MEMORY .....	4
SWITCHLESS INTERCONNECT .....	5
<i>Figure 1: System Interconnects - EV68 GS160 and EV7 Marvel.....</i>	<i>6</i>
SYSTEM CLOCK.....	7
I/O HARDWARE .....	8
POWER AND COOLING INFRASTRUCTURE .....	8
SOFTWARE FEATURES ENHANCING PLATFORM RAS .....	9
<b>FINAL THOUGHTS .....</b>	<b>9</b>





# *EV7 AlphaServers Deliver Enhanced RAS and Powerful Performance*

## **EXECUTIVE SUMMARY**

Launched in early 2003, the AlphaServer ES47, ES80, and GS1280 systems represent the latest generation of high-performance servers to satisfy the demanding requirements of Tru64 and OpenVMS customers. Known under the project code name “Marvel” while under development first at Digital Equipment, then Compaq, and now Hewlett-Packard (HP), these systems offer the top-notch performance that AlphaServer customers have come to expect. Every bit as important, these new servers feature enhanced RAS (Reliability, Availability, Serviceability) capabilities that position them to be the centerpiece of business-critical computing.

Historically, the blazing performance leadership of the Alpha processor proved irresistible to customers dealing with compute-intensive workloads, especially the high-performance technical computing (HPTC) community. The new EV7 systems do not disappoint this constituency – improved memory and I/O bandwidths enable the Alpha processor to deliver impressive performance. Particularly noteworthy, the new servers achieve powerful performance with an innovative system design that enhances their RAS characteristics.

By implementing switchless system interconnect, L2 cache, and memory controllers on the EV7 die, the AlphaServer developers dramatically reduced the number of components and connections that could contribute to failure. That architectural simplification is estimated to extend the average interval between failures by as much as 30%. Innovative features such as RAID memory and memory troller background correction help avoid uncorrectable memory errors even as the amount of memory grows very large. Error Correcting Codes (ECC) are widely used to protect data paths – HP indicates that over 90% of the EV7’s signal pins are covered by ECC or parity, and over 90% of the chip’s circuitry is protected by ECC. Extensive use of N+1 and hot-plug capabilities in the power and cooling subsystem helps to minimize failures of this underlying infrastructure. On top of the hardware features, enhancements in Tru64 UNIX and OpenVMS support multi-path I/O and predictive failure analysis.

## INTRODUCTION

Although certainly not immune to the effects of the economic environment, demand for computing continues to grow. In fact, computing plays an increasingly critical role for technical users and mainstream business customers. Engineering, scientific, and academic developers, and researchers look to take advantage of improved price/performance to attack larger, more complex problems. Enterprises seek a differentiated edge over their competitors, for example, through increased productivity or enhanced customer relationship applications. Additionally, consolidation of distributed servers into an efficiently managed, centralized infrastructure further drives the need for larger, more powerful systems.

Reflecting the ever-growing demand for computing power, the bragging rights of benchmark leadership often dominate the IT headlines. Performance is significant, for both HPTC and enterprise customers. At the same time, these critical computing resources must remain up and running to deliver that performance. Customers often take for granted that their systems will be robust and resilient. And indeed, those assumptions are well founded; today's platforms are highly reliable workhorses. However, such dependable servers do not "just happen." They result from a well-planned design that often remains unrecognized. This paper looks at some of the "behind-the-headlines" design decisions that assure the latest AlphaServers will meet increasing customer expectations for robust and resilient computing.

Let us begin by placing server failures into perspective. Customers recognize that most causes of unavailability can be traced to human operational errors. While the proportion of human-induced failures varies considerably, many customers feel about 60% of their failures can be attributed to human error. Some are just plain mistakes that should have been easily avoidable. Others involve confused decisions stemming from unclear or complex operational procedures. Part of the solution can involve automating some of the routine procedures. But because so much is *not* merely routine, detailed procedures and objectives need to be documented clearly and followed by comprehensive staff training. While vendors can guide customers in this process, improving operational errors lies with the customer.

The remaining 40% or so of failures typically splits fairly evenly between software and hardware. Many of the software failures relate to integration incompatibilities or change management issues. Since most customers manage somewhat unique workloads, resolution falls primarily on the customer. The customer needs a disciplined approach involving thorough stress testing of applications under conditions that represent production environments, as well as a clear process for regression testing of patches and fixes before applying to the production environment. Hardware failures represent the remaining 20%. While not a large number, for the most part these problem areas can be addressed by the hardware vendor, rather than by the end user.

In addition to designing for high performance, the Marvel development team was chartered to enhance the RAS characteristics of this new generation of AlphaServers. The developers examined predecessor AlphaServer systems to understand field failure scenarios, and reviewed failure rate statistics of the underlying components.

### **Performance of the New AlphaServers**

Announced in January 2003, the 1.0 GHz EV7 Marvel midrange includes the up-to-four-way ES47 and the up-to-eight-way ES80. At the high end, the 1.15 GHz GS1280 currently ships in eight- and sixteen-processor configurations. The GS1280 is the new flagship of the AlphaServer Line.

Building on a long-standing reputation for high performance, the EV7-powered AlphaServers can be expected to deliver competitive performance even in the face of increased pressure from other recently updated chips. Benchmarks have been reported for the new AlphaServer systems that illustrate performance leadership across a broad range of applications. Compared to predecessor EV68-based systems, the EV7 servers are projected to offer 35% to 50% or greater performance under Tru64 UNIX. OpenVMS users are expected to see even greater performance gains. In general, larger configurations benefit more from the high system and memory bandwidth coupled with low access latency of the mesh switchless interconnect. Thus, larger GS1280 configurations deliver even more performance relative to other systems.

On the Oracle 11i Application Standard Benchmark, an eight-way GS1280 measured 7728 users, which leads all other eight-way systems. The GS1280's SPECint\_rate/SPECfp\_rate of 536/313 currently leads all other measured 32-processor systems. In the STREAM memory bandwidth benchmark, the GS1280 delivers the best memory bandwidth for non-vector 16-processor systems, and far exceeds the memory bandwidth performance of any other non-supercomputer RISC-based systems with configurations containing twice as many processors. On the SAP SD 2-Tier Standard Application Benchmark, a 32-way GS1280 beat all other 32-way systems. Reported results for OpenVMS customers in telecommunications and healthcare show performance gains of 100% over previous-generation AlphaServer systems.

In the ever-leapfrogging benchmarking race, others may slip ahead of the EV7 AlphaServer systems at some point. One point is clear – the new AlphaServers yield impressive memory and I/O bandwidths as well as lower latencies; characteristics that will deliver strong performance for both technical and commercial computing environments.

The resulting EV7-based AlphaServers rely on a simple, straightforward architecture that reduces chip count to inherently improve reliability. Supporting infrastructure, such as power supplies and fans, have also been redesigned to improve RAS. AlphaServer developers estimate that Mean Time Between Failure (MTBF) has been improved by 15% to 30% thanks to the simplified architecture. Furthermore, single system availability has been extended, since virtually all hardware failures are backed up by redundant components. In addition, many of the components with the lowest MTBF – power supplies, fans, and management system service processors – can be hot-swapped while the system continues to run. Details of the enhancements form the subject of this paper.

Even with increased efforts to reduce outages, unexpected failures caused by hardware, software, or personnel will occasionally result in unplanned downtime. Failover clustering remains an important option for continued application availability. Single-system image clustering, as found in TruCluster Server from HP, dramatically simplifies cluster management. Reducing the complexity of the cluster by managing it as a single system, the opportunities for human error are greatly reduced. Note however that clustering is not true fault tolerance, and application users may still suffer some disruption before the backup system takes over fully. Improved single-system RAS can reduce the occurrences when failover clustering is called upon. That is, enhanced single-system RAS complements failover clustering. The combination of a robust and resilient single system and failover clustering provides a solid, dependable computing environment.

## **KEY RAS IMPROVEMENTS**

EV68 Alpha chips use 15.2 million circuits; the EV7 chips contain over 150 million circuits. Certainly a primary goal of the Alpha designers was to employ the additional circuits to drive performance higher. At the same time, using those circuits in a way that simplified overall system design to deliver a highly available, dependable platform stood out as a key issue. The descriptions below highlight how the AlphaServer team enhanced RAS at the same time it created the high-performance server sought by its customers. While not an exhaustive list of all RAS features, the breadth of areas described indicates the concerted drive to create a powerful and reliable server.

### **ON-CHIP L2**

The EV68 processing core remains a highly respected, high-performance compute engine. Alpha designers wished to preserve that core so that applications would not need the re-optimization that typically accompanies the introduction of a new core design. Since the processor core was not limiting performance, the Alpha team addressed the task of feeding data to the core fast enough, instead of implementing a new core. EV68 employed an off-chip Level 2 (L2) cache, a common design choice when the EV6 was first designed. Exploiting the larger circuit capacity of the EV7, Alpha designers brought L2 on-chip, substantially boosting performance, thanks to the low latencies of an on-chip cache. Equally important, the elimination of off-chip L2 memory chips, chip carriers, and sockets dramatically reduces the number of parts and interconnections with a corresponding reduction in possible failures. As with the L1 caches, the on-chip L2 is covered by ECC that corrects single-bit failures and detects double bit failures.

### **RAID MEMORY**

The Alpha designers also brought the control circuitry for addressing main memory onto the EV7 chip. Once again performance improves: the on-chip memory controllers allow much lower memory access times and higher memory bandwidths than an external memory controller. In addition, reliability is enhanced: the elimination of external chips and interconnections also removes potential points of failure. Furthermore, the EV7 team added another layer of hardware error correction beyond ECC - a RAID memory option.

In most servers today, memory is protected by means of a Single Error Correct, Double Error Detect (SECDED) ECC scheme. Memory error rates are fairly low and SECDED is usually adequate since it is reasonably rare that more than one error will occur in the same word of memory. However, high-end servers can have vast amounts of memory installed, hundreds of gigabytes today, soon stretching into the range of 512 GB - 1 TB. Although the failure of an individual memory location may be tiny, when trillions of bits are considered, the chance of a double-bit error somewhere raises a concern.

Most memory errors are soft errors, the inadvertent loss of a data bit due to electrical interference, including cosmic rays. The memory circuits themselves remain fully functional and can be reused to hold data. In contrast, a hard error signifies permanent failure of the memory circuitry for that bit. If a soft error can be corrected before another error might occur in the same memory word, most double-bit errors can be avoided. Tru64 UNIX features a memory “troller” mechanism that proactively scans AlphaServer memory for single bit errors, allowing the hardware to automatically correct them using ECC. By cleaning up these intermittent soft errors, the troller significantly reduces the probability that two errors will eventually appear in the same memory word. By keeping track of where single bit errors were located, the troller also can identify memory locations suffering repeated failures, and allows the operating system to avoid using that memory before it turns into a permanent, hard failure.

Once in a while, hard errors do occur, and may even disable multiple memory bits. For those occasions, something beyond SECDED ECC or memory troller is called for. To address such problems, the EV7 on-chip memory controller incorporates a provision for RAID memory.

Similar to a RAID disk configuration, the RAID memory option adds a redundant-memory RIMM (RDRAM In-line Memory Module) that can be used to correct for multiple failures within one of the other RIMMs. The base EV7 AlphaServer memory configuration spreads data across four RIMMs per memory port. RAID memory adds a fifth RIMM that can be used to recreate the data from any one of the other four RIMMs. The memory controller uses the standard memory ECC to identify which RIMM has failed, and passes reconstructed, correct data along to the processing core. Since the entire bad RIMM is being bypassed, maintenance can be scheduled at a time convenient to the customer’s operational and service-level agreement requirements.

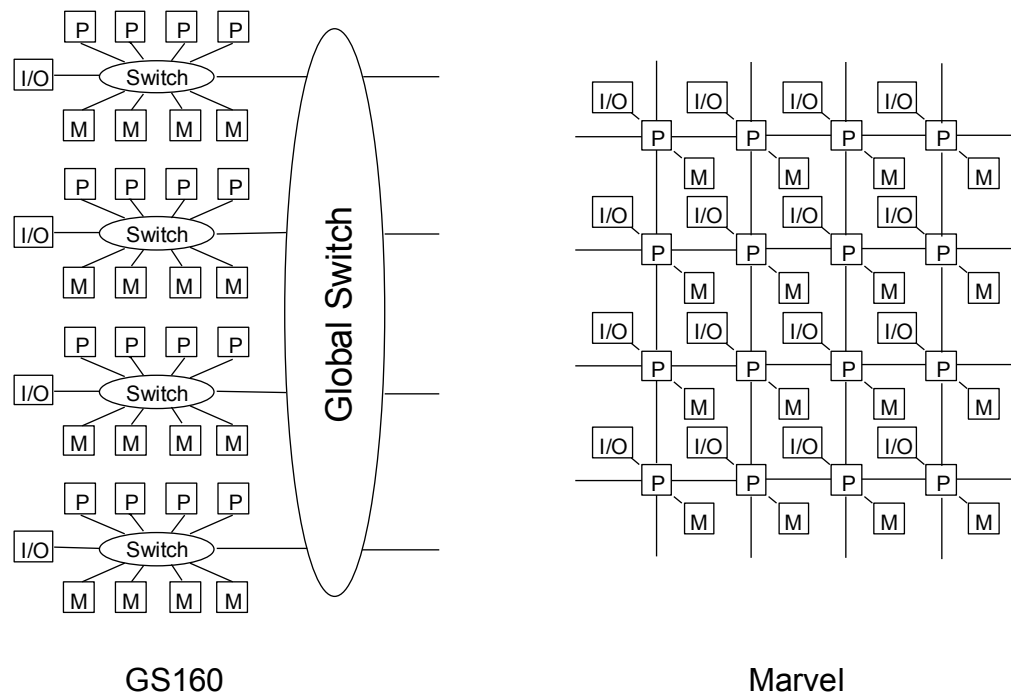
## **SWITCHLESS INTERCONNECT**

A major challenge to system developers is to design a scaleable interconnect that efficiently extends to large configurations but that does not penalize smaller systems with the costs of an elaborate interconnect infrastructure. Marvel designers addressed this challenge by integrating the processor interconnect within the EV7 chip. Each EV7 chip contains an internal router that directly connects the processor core to local memory, a high-performance I/O bus, and four connections to the routers on adjacent EV7 chips; all as passive interconnect without the need for additional external logic chips.

Figure 1 illustrates sixteen-processor configurations of the EV68 GS160 system, and the EV7 “Marvel” GS1280. As the left side of the figure shows, the EV68 systems employed a hierarchy of switches to interconnect processors, memory, and I/O. The right side highlights the Marvel design, which does not require any external active electronic switches for system interconnect. Every EV7 processor brings its own portion of the interconnect. A four-processor Marvel configuration (one row of processors in the figure) could easily expand to 64-

processor systems by adding rows and columns of processors to the mesh. (To be accurate, the left and right edges of the mesh connect to each other, as do the top and bottom edges. Thus, with the edges wrapped around, the EV7 interconnect is properly called a 2D torus topology, or donut in more conventional terms.) Switch-based systems typically cover a narrower range of configuration sizes since switches cannot easily be expanded in small increments. For example, the GS160 illustrated in Figure 1 carries the cost of the full Global Switch needed to support a 32-processor GS320. 64-processor EV68 systems are not offered since a switch to interconnect twice as many processors would require far more circuitry.

**FIGURE 1:**  
 System Interconnects -  
 EV68 GS160 and EV7  
 Marvel



Not only does the EV7 internal router permit efficient scalability from small-to-large configurations, it also improves performance, since the internal router circuitry operates faster than the external switches. And, the elimination of local and global switches vastly reduces the number of chips and connections that carry inherent failure rates.

The Marvel interconnect is protected by SECDED ECC to recover from intermittent bit failures. The ECC correction is applied “on the fly,” so as not to affect overall system interconnect bandwidth or latency.

ECC is intended to overcome individual bit errors and cannot correct for permanent interconnect failures. Since the GS160/320 switches did not contain redundant paths, a switch failure could isolate processors and memory, necessitating a switch repair to regain full use of the system. Even hard partitions on EV68 systems could all fail together, if the switch failed. On the other hand,

failure of an interconnect link in Marvel systems may cause a crash in the affected hard partition, but nowhere else in the system. Upon reboot, Marvel's 2D torus automatically reroutes around defects.

As part of normal operation, Marvel's interconnect network monitors traffic on each of the links to detect congestion hot spots. If the network seems overly busy along a particular path because of a large number of data packets, subsequent packets of data will be adaptively routed to their final destination along an alternate set of interconnect paths. An actual router failure will be detected by the same network traffic monitor, which will then reroute around the failure during reboot.<sup>1</sup>

The failure of the computing core of a processor does not require the router portion of the EV7 to be taken out of service. Rather, the core processor functionality can be disabled while the memory, I/O connections, and interconnect links remain in use. In that way, other processors in the system can continue to access that memory and I/O, and the network can go on routing packets along the paths passing through that chip. Similarly, a failure in the I/O access circuitry can be isolated to permit memory and interconnect to remain accessible.<sup>2</sup>

## **SYSTEM CLOCK**

Typical large systems require elaborate clock distribution circuitry to guarantee that clock signals are precisely aligned across the entire system. The interconnect mechanism of the EV7 router function has been designed to tolerate misalignment of clock signals between the EV7 chips. (In technical terms, the EV7 uses pseudo-synchronous clocking to transfer data between chips using clock forwarding.) While clock distribution failures are uncommon, the use of clock forwarding illustrates that the EV7 team devised simple, straightforward designs. They did not rely on complex techniques that might carry higher failure rates.

In essence, each EV7 processor incorporates its own clock that supports the EV7 processor, interconnect, memory, and I/O. The predecessor EV68 systems used a single clock for the entire system, albeit implemented with extremely reliable components. Because the GS160/GS320 clocking used very low failure rate parts, with a mean time between failure calculated in decades, there were no customer complaints regarding the clocking. Nonetheless, for Marvel, the AlphaServer took the extra step of replicating clocks for simplicity of design and enhanced availability.

---

<sup>1</sup> When the system reboots, firmware performs an interconnect integrity test. If that test fails then the system will map around the failed link. If the interconnect integrity test passes, then the link is used.

<sup>2</sup> If the system reboots and the core fails self-test, then its memory and I/O will be unavailable. If the core fails after the system is up and running, and the indictment software determines it is safe to off-line the core and leave the system running, then the I/O and memory will be accessible. If the core fails in a more catastrophic manner, then it will crash the system, and more than likely all of its memory and I/O will not be accessible on reboot.

## I/O HARDWARE

The EV7 chip was not alone in taking advantage of large circuit count to absorb functions previously contained on external chips. The I/O chip “IO7” ASIC (Application-Specific Integrated Circuit) integrates the functionality previously contained on eight separate chips in the GS160/GS320 implementation. Marvel also added the functionality of AGP support in this chip, which the EV68 systems did not implement. Once again, fewer chips and fewer chip interconnections can improve reliability dramatically.

Each IO7 chip drives a set of paths to remote I/O drawers containing PCI/PCI-X and AGP slots. ECC protects all I/O data paths within the drawers as well as those connecting to IO7 chips. The I/O drawers offer hot-plug capability so the I/O adapter cards can be repaired/replaced without taking the system down.

## POWER AND COOLING INFRASTRUCTURE

As discussed above, the scaleable interconnect of the EV7 allows the same design point to serve low-end through high-end configurations. While the electronics remain common across systems, there is an added advantage: the power distribution and heat removal technologies can also be common across the product line.

Power supplies, voltage regulators, and fans typically carry higher failure rates than the logic circuitry. Specifying components to achieve high reliability at reasonable cost involves a tricky set of tradeoffs. (Money spent on power and cooling does not increase the benchmark performance and only raises overall system price.) Deploying a common power/cooling technology across low- to high-end systems can increase procurement volumes enough to drive down costs even when specifying premium components. For example, AlphaServer developers indicate that a common design point allowed them to employ a more reliable power supply, at an attractive price. By comparison with the more standard selection of different power supplies for each model in the product line, AlphaServer chose the cost-effective route.

All Marvel systems, from two processors to sixty-four processors, possess redundant or N+1 power supplies, voltage regulator modules (VRM), and fans. Should one of these infrastructure components fail, the system continues to run. Except for VRMs, hot-swap capabilities allow the failed unit to be removed and replaced while the system continues to operate. Dual AC input allows configuring a second AC power source as backup.

In a similar vein to the single/dual clock design change, the EV68 systems employed a single high-performance blower with a calculated mean time between failures that far exceeded other server components. However, even a minuscule chance of failure concerned some customers since the GS160/GS320's blower was not duplexed. Now, with a redundant, hot-plug fan assembly used across all Marvel systems, a cost-effective alternative removes any single point of failure concern.

## **SOFTWARE FEATURES ENHANCING PLATFORM RAS**

The Tru64 memory troller was highlighted earlier. Most of the discussion so far has focused on hardware implementations specific to the Marvel platforms. There are additional software features, in both Server Management software and in the operating system, that complement the hardware to further enhance system RAS.

As mentioned earlier, a failure of part of the system interconnect 2D torus can be overcome by adaptive rerouting. Thus, I/O circuitry within an EV7 should not become isolated due to a system-interconnect failure. But if a critical I/O device is connected to only one IO7 controller, the device can be lost due to a failure in that controller. The preferred alternative is to configure multiple paths to critical I/O devices, each from a different EV7 I/O controller. Then, in case of failure, the Tru64 or OpenVMS operating system can reach the device through its Dynamic Multi-path I/O support.

In conjunction with the server management hardware embedded within the systems, Tru64 5.1B can monitor resources and track intermittent failures of processors, memory, disks, and various I/O adapters and devices. Recognizing that repeated failures may indicate an impending permanent, non-recoverable failure, the operating system can suspend further use of the suspect resource and log it for deferred diagnosis and/or maintenance.

## **FINAL THOUGHTS**

Historically, Alpha processors were renowned for their fast clock rates delivering blazing floating-point performance. But AlphaServer developers understand that high performance derives not just from clock frequency or complex superscalar design; it also requires a robust cache/memory system to feed the insatiable appetite of high-performance processors. For Marvel, the AlphaServer design team focused on using EV7 circuits to provide high bandwidth and low latency paths to memory, I/O, and other processors in the system. In addition to boosting performance, the simple design enhanced RAS as well by eliminating failure-prone chips and connections. In addition, N+1, hot plug, power, and cooling further enhance dependability of the systems. This double win – higher performance and enhanced RAS – positions the EV7 AlphaServer as an attractive platform for existing Tru64 and OpenVMS customers.

In 2004, a fabrication shrink will allow the EV79 to achieve faster clock rates and corresponding higher performance. Since the EV79 processors can be added to EV7 systems, investing in Marvel today allows additional growth next year.

In the longer term, the AlphaServer road map entails an evolution to Itanium-based systems running HP-UX or OpenVMS. AlphaServer users should be planning their testing, pilot, and rollout scenarios and can begin that transition at any time. While such testing and piloting are underway, those customers want to ensure that their production workloads continue running on solid, dependable, high-performance platforms. The new EV7 AlphaServers are those platforms.