

Resilient Systems, Inc.
199 Nathan Lane
Carlisle, MA
01741 U.S.A.
(tel) 1.978.369.5356
(fax) 1.978.371.9065



WWW.RESILIENSTYS.COM

White Paper

July 2003

VAX Emulator on HP's Marvel AlphaServers Extends the Life of Legacy DEC VAX Systems

Combination of **GS1280** and **CHARON-VAX/AXP Plus**
Provides Superb Server Consolidation for VAX Installed Base

**Marvel Partitions and CHARON-VAX Cluster Capabilities Permit VAX
and Alpha Systems to Form Mixed Clusters within a Single System Cabinet**



Strong synergies in the latest technologies from Hewlett-Packard and Software Resources International promise not just a reprieve for remaining VAX systems, but a clear path to 21st century platforms. VAX/VMS users are now able to take advantage of ever-faster chip speeds and massive storage platforms like millions of users of modern technology.

This whitepaper describes the results of rigorous testing performed by Resilient Systems at Hewlett-Packard's Littleton, Massachusetts lab, using hardware and test suites provided by VMS Engineering. The system provided for this testing was a 16-way GS1280 AlphaServer (code-named Marvel) running multiple copies of CHARON-VAX/AXP *Plus*. The test suites were the same used in previous decades by VMS Engineering to test new VAX hardware designs.

The results show that the combined strengths of these products permit **server consolidation** and **single-platform clusters** while providing enhanced performance in one Alpha footprint.


The Marvel of It All

Based on the new EV7 processor, the GS1280 AlphaServer delivers an unprecedented combination of performance, scalability, and system reliability. The architectural advancements of Marvel over previous switch-based NUMA systems like the 32-way GS320 (code-named Wildfire) are numerous. All elements required for symmetric multiprocessing now reside on a single chip. In addition to an on-chip L2 cache, two on-chip memory controllers provide exceptional memory bandwidth. In an industry-first achievement, an on-chip router connects AlphaServer processors directly to one another. This "switch-less" mesh design results in very high interconnect bandwidth up to 64 CPUs. SPEC_rate 2000 tests on a 32-way Marvel proved the GS1280 achieved nearly 100% linear scalability.

The I/O performance and scaling of Marvel versus Wildfire is equally impressive. The GS1280 provides flexibility in configuring I/O, from one I/O chip per system to one I/O chip per processor. The result is a platform with linear scaling in I/O, yielding eight times the I/O bandwidth of the GS320. Moreover, Marvel's Lego™ block design of hot-swappable components results in a robust platform with 15 to 30 percent improvement in Mean Time Between Failure (MTBF) over the previous Wildfire generation. Available in multiple N-way configurations, the enterprise-scale AlphaServer GS1280, along with departmental and workgroup models ES80 and ES47, provides significant performance and reliability improvements over previous generation GS320 and ES45 systems.

CHARON-VAX/AXP *Plus* the Benefits of Alpha

The new VAX-on-Alpha emulator from Software Resources International takes full advantage of these evolutionary improvements of Marvel. Software Resources International specializes in the migration of operating systems and applications to modern platforms (e.g. OpenVMS Alpha to Itanium), and the development of hardware emulators for PDP and VAX. The emulators are mathematical models of the hardware architecture; written in C, they run as ordinary applications on modern platforms, as **Figure 1** illustrates.



No code conversion is required when using CHARON-VAX/AXP *Plus*. Simply use backup/image to transfer existing VMS and application binaries to Alpha.

CHARON-VAX/AXP *Plus* is the second-generation of Software Resources International's VAX hardware emulator for Alpha. The first emulator modeled a MicroVAX 3600 in software. The new emulator provides the functionality of a VAX 3100 Model 98 hardware system, complete with up to 512MB memory, dual SCSI storage buses, and a 10/100 Mbps Ethernet network.

Combined with Marvel's scalability and reliability, the sophisticated instruction preprocessing now available in CHARON-VAX/AXP *Plus* has significantly increased the viability of preserving business-critical VAX applications via VAX emulation. One or many low to mid-range MicroVaxen can be replaced by entry-level ES47s. And testing has now shown it is even possible to replace one or more VAX 77xx or VAX 9000s – the very high end of the VAX range - on an N-way GS1280.

Proven Performance

Specifically, the testing conducted at HP's Littleton, Massachusetts laboratory in Spring, 2003 by independent research firm Resilient Systems, Inc., proved that a 16-way Marvel GS1280 running CHARON-VAX/AXP *Plus* delivers the equivalent of a VAX 3198 or VAX 7610 – over 36 VUPs – on each CPU of an AlphaServer system.

Even more impressive, the remarkably efficient CHARON-VAX kernel (0.5 MB) achieved the same scalability as the underlying Marvel hardware when running multiple instances of CHARON-VAX/AXP *Plus*. As the following graph shows, the compute power – measured in traditional VAX Units of Performance (VUPs) - obtained by running additional instances of the VAX emulator scaled nearly linearly as the number of processors increased.

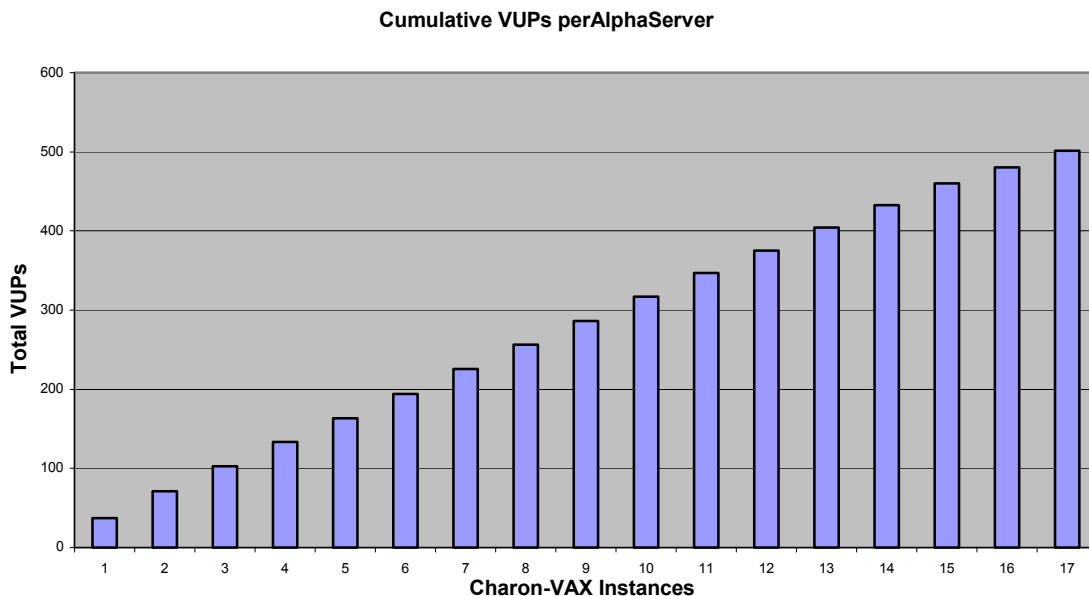
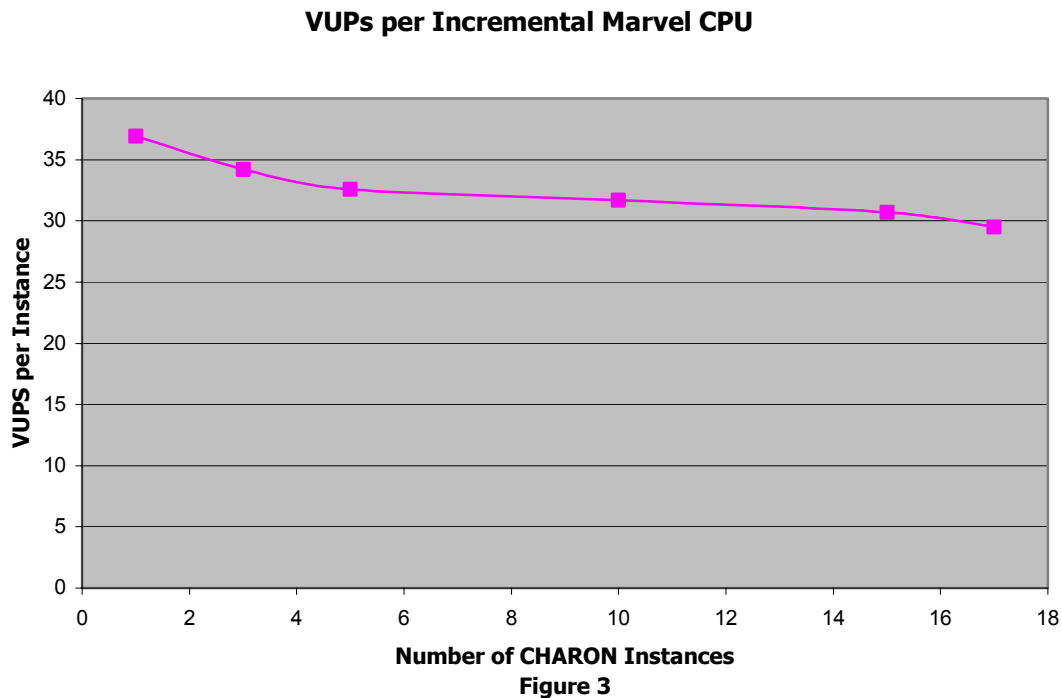


Figure 2

The software architecture of the CHARON-VAX emulator consists of two threads – one thread to execute the emulator and a second thread to field interrupts, run the scheduler, manage resources, handle I/O to storage devices and manage network I/O. While it is possible to run

both threads on the same processor, for optimum performance the emulator thread should have 100% of a CPU available to it. The second thread, automatically assigned to a separate CPU when one is available, requires a fraction of the compute power available to it. The remarkable synergy between the hardware architecture of Marvel and the software architecture of CHARON-VAX produced an optimum configuration of 15 instances of the emulator on a 16-way AlphaServer, with the 16th CPU managing resources for the other 15. Only when the number of instances of CHARON-VAX exceeded the number of CPUs was significant contention observable.

Specifically, Resilient Systems' tests proved each instance of CHARON-VAX/AXP *Plus* delivers an average of 32 VUPs on an AlphaServer with 16 CPUs, each independently running an instance of the software. As the graph below demonstrates, performance only declined below 30 VUPs per CPU when the number of CHARON-VAX executables exceeded the physical number of CPUs.



The Benefits of Consolidated Power

The results of these tests clearly indicate that multiple individual VAX servers or VAXstations could easily be consolidated on the same Marvel host. **Server consolidation** offers many benefits to VAX sites with multiple remaining systems, including reduced footprint and power consumption, and *greatly reduced* hardware maintenance costs. Dependence on increasingly scarce VAX/VMS system engineers is lessened by reducing overall staff requirements, and the risk of business disruption due to malfunction of aging hardware is alleviated by the superb reliability and MTBF of Marvel.

Similarly, **single platform clusters** can now be created - an entire cluster of existing VAXen could be recreated as multiple cluster members of the same cluster on a single Marvel host. Or the configuration could be aggregated and then spread over redundant Marvel systems to attain the highest possible availability through the independence of separate hardware systems. The benefits would include all the benefits of server consolidation described above, *plus* the failover capability inherent in VMS clusters.

Please note that in either scenario – server consolidation or single platform clusters - VUP performance numbers of existing VAXen must be carefully totaled to not exceed the maximum recommendations for CHARON instances on Marvel. The following chart illustrates the cumulative capacity – in VUPs – of the various Marvel platforms.

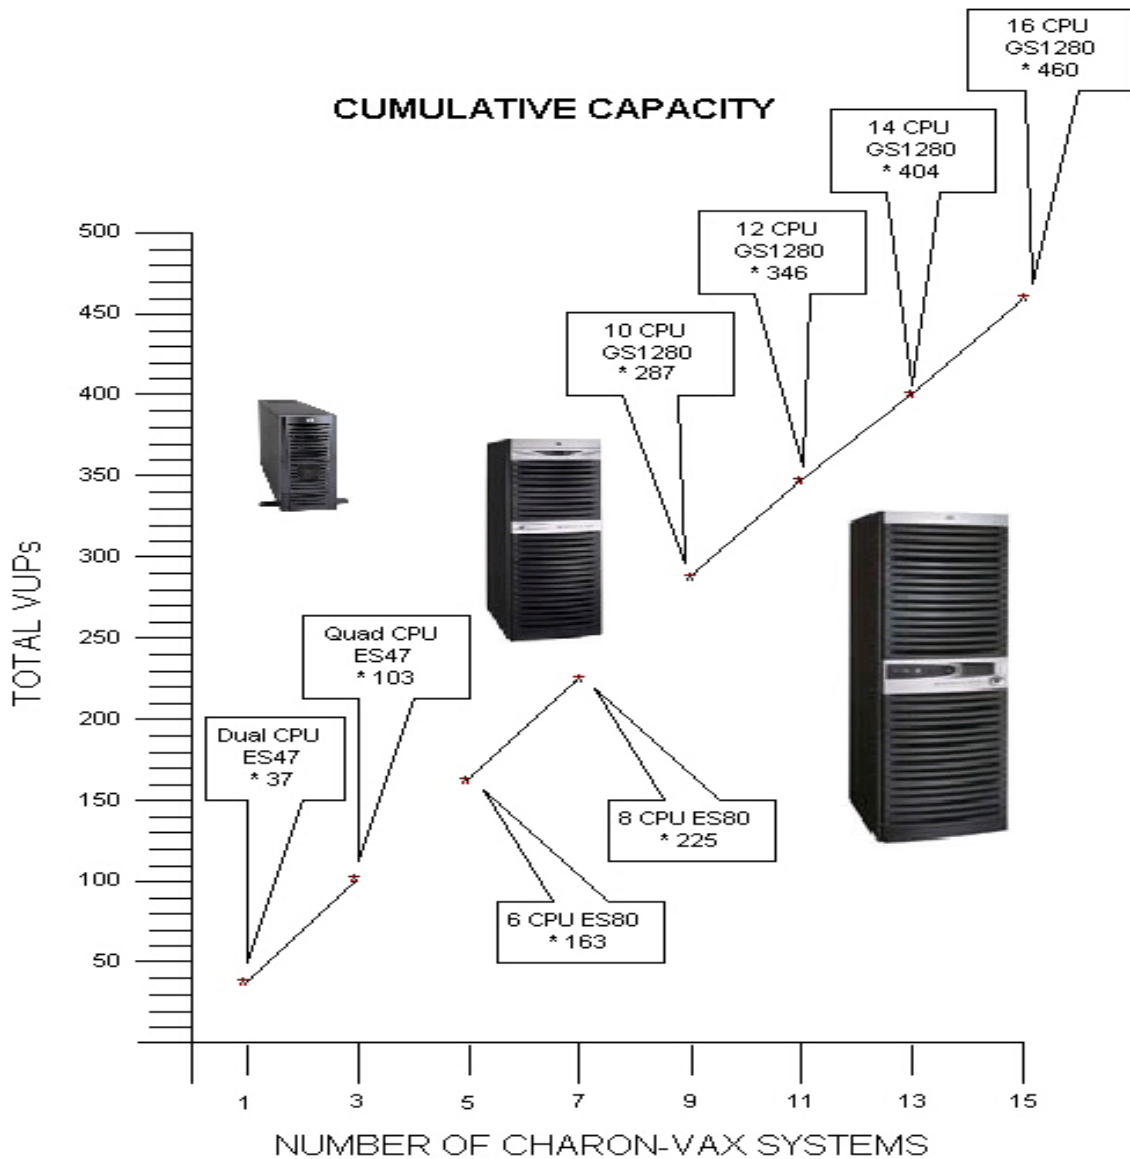


Figure 4

I/O Capacity Keeps Pace

Furthering the synergy between the Marvel hardware and the CHARON-VAX software, Resilient Systems' tests have shown that the bandwidth available to the emulator is nearly identical to what is physically attached to the AlphaServer host. Repeated testing showed that native Alpha disk transfers achieved 4.47 MB/sec when accessing a local SCSI disk versus 4.45 MB/sec for the CHARON-VAX emulator when accessing the same physical disk. In other words, emulator overhead is less than 1% for tasks such as disk-to-disk file copy operations or VMS backup transfers.

This testing has proven that customers are now able to assimilate high performance storage subsystems, such as Fibre Channel, into a legacy VAX configuration. The ability to increase storage capacity by transforming VAX physical disks into 'disk image' files on the replacement platform is inherent in CHARON-VAX. When combined with the AlphaServer's support for robust storage technologies, critical VAX applications can now take advantage of storage capacity and I/O throughput unimaginable in the heyday of the VAX.

The network device is key to integrating an instance of the CHARON-VAX emulator with other DECnet nodes, other cluster members, or, via IP, with corporate LANs and WANs. This channel also provides user connectivity through telnet and third-party terminal emulators, and thus is a key component of a VAX replacement configuration. When the devices were set to match the 10baseT adapter of a VAX system, Resilient Systems observed data rates through the network device at over 1.8Mb/sec for sustained data transfer, and near the full 10Mb/sec possible for messages. 100 Mbps Ethernet adapters can be used with the current version of CHARON-VAX/AXP *Plus* but were not tested. For operation at 100 Mbps, an Alpha SMP host with a CPU frequency of at least 1 GHz is required. Network throughput can be individually tuned for specific protocol classes (e.g. DECnet, TCP/IP or VMS cluster communication).

Mixed Architecture Clusters in a Single Marvel Box

In May 2003, HP announced support for hardware partitions on Marvel. The ES47, ES80 and GS1280 support this option, and CHARON-VAX/AXP *Plus* is likewise able to take full advantage.

Hardware partitions permit multiple instances of the OpenVMS operating system to run concurrently in physically separate parts of the computer. Such a configuration facilitates the dedication of partitions to specific applications, with the ability to tune and secure each partition to the specific demands of its application set. By effecting the partition of the system into multiple independent Alpha processors, this new feature facilitates the deployment and execution of multiple instances of CHARON-VAX/AXP *Plus*. CHARON-VAX can be run as an application on one or more of the CPUs in a processor partition, or across multiple CPUs in multiple partitions.

This permits the construction of a variety of mixed VAX and Alpha configurations, all within a single system cabinet. See **Figures 5 and 6** for some examples.

MIXED VAX and ALPHASERVER CONSOLIDATION EXAMPLE




Figure 5

For example, the 8-way AlphaServer in the drawing above could be partitioned into three separate Alpha systems. The Alpha system with the five CPUs (in yellow) could be configured to run multiple CHARON-VAX instances.

Alternately, an identical 8-way AlphaServer with three partitions could be configured to create a 6-node, mixed architecture cluster, as in the example below.

MIXED VAX and ALPHASERVER CLUSTER EXAMPLE




Figure 6

In this view, three of the eight available CPUs would run as actual Alpha nodes (one single CPU node and one dual-CPU multi-processor node). The remaining five CPUs would run four instances of CHARON-VAX/AXP *Plus*, with the fifth CPU fielding user interrupts and managing disk and network I/O as previously described.

The Test Suites

In addition to providing Resilient Systems with access to the Marvel configuration and lab facilities to conduct these tests, VMS Engineering provided the diagnostic suite used in prior years to test new VAX hardware designs. The tests verify conformance of the new hardware to expected test results to ensure proper execution of the VAX instruction set. The comprehensive suite exercises nearly every VAX instruction, including all 3-operand VAX instructions as well as single, double, and floating point calculation speeds. For some instructions, the CHARON-VAX emulator was more than 10 times faster than any real VAX.

In addition to this suite, Resilient Systems used the VUPs Calculator utility, which is a mix of fixed and floating point instructions, to test individual CPU performance. To test scalability of CHARON-VAX on multiprocessor configurations, Resilient used standard Dhrystone tests because they produce much more granular results than the VUPs Calculator. Resilient first determined peak Dhrystones on a single CPU, and then ran simultaneous Dhrystone tests on multiple instances of CHARON-VAX in multiple N-way Marvel configurations up to a maximum of 18 instances on a 16-way Marvel. These tests were repeated over three days and the average Dhrystone performance was calculated. The standard formula to convert Dhrystones to VUPs was then used to produce the graphs in Figure 2 and Figure 3.

The complete test results from the VMS Engineering test suite and results of the CPU, disk, and network performance tests and CHARON-VAX scalability tests are available from Resilient Systems. Please contact Nancy Lyons, President, E-mail Address n.lyons@resilientsys.com.

Summary

The remarkable synergy in performance and reliability of HP's new Marvel AlphaServers and Software Resources International's new *Plus* version of CHARON-VAX/AXP offer VAX/VMS customers a rock-solid platform for server and cluster consolidation and an economical means to preserve vital VAX applications for years to come at greatly reduced maintenance cost and business risk.

For more information about Resilient Systems, Inc., go to

www.resilientsys.com

For more information about Hewlett-Packard and Marvel AlphaServers, go to

<http://h18002.www1.hp.com/alphaserver/>

For more information about Software Resources International S.A. go to

www.softresint.com