THE CLIPPER GROUP
Navigator™
*Navigating Information Technology Horizons*SM

Published Since 1993 Report #TCG2009004 January 23, 2009

Virtualizing Open Systems Servers with SMP — ScaleMP Lowers TCO, Raises Performance

Analyst: David Reine

Management Summary

In 2008, consumers in the U.S. saw an unprecedented rise in the cost of gasoline. In some parts of this country, gasoline rose to over $5.00 per gallon. We received no sympathy from European drivers, however, as they looked at gas prices in the U.S. as a bargain. This does not even address the impact of fuel costs on the homeowner, trying to heat their houses to ward off frostbite in a bone-chilling winter. Even with the subsequent drop in gas prices to under $2.00 per gallon, the task of reducing fuel consumption for major enterprises and governments everywhere has moved to a crisis status as budgets become frozen – not by the elements, but by a spine-chilling recession. The *obvious* answer for the manager of any fleet of vehicles is to replace the existing fleet of sedans with smaller cars that deliver better fuel economy. In Europe, for example, with gasoline priced at an equivalent of well over $5.00 per gallon, the *Smart Car,* with a mileage rating of better than 40 MPG, has become popular. The Smart Car, known as *smartfortwo* for a reason, will only seat two people. What does the transportation manager do if he has to move 10 or 12 people between the same two points? Would it make more sense to buy a single Minibus with an MPG rating of 10 to 15 MPG, but seating 12, 15, or more? Driving six Smart Cars to transport 12 people to the same destination 40 miles away costs more, burns more fuel, and decreases air quality. At 40 MPG, the Smart Cars would burn 6 gallons of gas while the Minibus would only consume 4 gallons (at the low end of their range). Under these conditions, driving a Minibus would save 33% on fuel costs, not to mention the savings on insurance, maintenance, and administration. **Sometimes the obvious answer is not so obvious!**

Solving the total cost of ownership (TCO) issues in the data center may not be as obvious, either! Uncontrolled server sprawl is leading to a financial disaster in the enterprise IT arena. Low-cost, open systems servers are multiplying like rabbits in all application environments, with a new server being deployed for every new application installed. As with the Smart Car, a scale-out architecture is not always the best solution. Deploying hundreds (thousands?) of rack-mounted or bladed servers can create an unmanageable tangle of complexity, leading to an administrative nightmare trying to manage hundreds of copies of *Windows* or *Linux*, not to mention a myriad of mission- and business-critical applications. In addition, the cost of replicating and patching all of this software across your scale-out environment also will increase your TCO.

On the other hand, scale-up architectures (think Minibus) have been around for decades – mainframe anyone? Today, however, you can scale-up with the same open servers used in a scale-out environment. With the availability of *vSMP Foundation* from ScaleMP, the enterprise now has the capability to virtualize high-end computing to deliver higher performance and lower TCO by aggregating multiple processors into a virtual SMP environment. To learn more about vSMP Foundation, please read on.

IN THIS ISSUE

Current Data Center Architecture

In order to take advantage of previously under-utilized CPU resources and to consolidate multiple legacy servers onto fewer platforms, the IT staff in many data centers has transitioned older servers with single core processors onto the latest platforms with multi-socketed, multi-core, multi-stream processors from Intel (*Xeon*) or AMD (*Opteron*). With even more CPU resources available to the data center, the typical IT staff is deploying a virtualized environment, sub-dividing each core into multiple partitions using a virtualization hypervisor, such as *VMware* from EMC or *Hyper-V* from Microsoft, among others, into a scale-out environment with hundreds or even thousands of physical servers deployed. In this way, enterprise applications requiring only a fraction of the processing capability of the server can share those resources with other applications, lowering the operational costs of running the application server and thus lowering the TCO. This is an ideal architecture for many file/print/web services applications.

However, what about the high-performance computing (HPC) applications that are running the mission-critical aspects of the enterprise? HPC capabilities are mandatory for enterprises wishing to accelerate and optimize innovation to drive greater success in order to compete in a demanding global economy, especially in times of a spiraling economic environment. Clustering, in the form of a scale-out architecture, has helped to lower the TCO of HPC for parallel, mostly compute-intensive, workloads, but it is still complex, and applications with large memory requirements are still very expensive. Some of these applications require more processing time than a single core, or even a single quad-core processor can provide. In many large enterprises, these applications have typically been deployed on a mainframe or on a high-performance RISC[1] processor running a variant of the UNIX operating System, such as *AIX* from IBM, *HP-UX* from HP, or *Solaris* on *SPARC* from Sun, all of which can be classified as *symmetrical multi-processing systems*, or *SMPs*, on a scale-up platform. Unfortunately, these high-end solutions typically carry a price tag beyond the reach of a mid-sized enterprise and many departments within a larger enterprise. These businesses require the same high-performance computing capabilities of their big brothers, but their budgets limit them to the range of systems referred to as open platforms designed for a commodity x86 architecture based upon dual-processor Xeon or Opteron configurations.

The typical mid-sized enterprise also needs to access a superset of the resources available from any single x86 system, even a high-end quad-processor available from many of the Tier 1 server vendors. They need to aggregate both their hardware and software resources into a larger virtual SMP environment, changing the existing paradigm of virtualization on an off-the-shelf x86 server environment. Where one could consider virtualizing in a scale-out environment as breaking up the application to fit the hardware, virtualizing in a scale-up environment is more a combining of hardware to satisfy the application.

Scale-Up vs. Scale-Out

Today, when deploying a scale-up SMP solution on a dual-socket x86 server, the data center is limited in the amount of processing power, memory, and I/O resources that any one platform can throw at an HPC solution. Some enterprises have tried to upgrade a single two-CPU server with a quad-socket platform.

Unfortunately, these platforms do not have linear scalability, so they cannot provide a 2x improvement in performance. Furthermore, on a per socket basis, quad-processor systems are more expensive, deliver lower compute density, and consume more power than a dual-socket server. Also unfortunately, proprietary RISC SMP systems such as those from HP, IBM, and Sun, while easier to manage from a CPU and I/O perspective, are even more expensive on a per socket basis and choosing them generally limits the data center to a single supplier.

Trying to deploy a scale-out cluster solution as an alternative to the SMP environment is unlikely to meet with success or management approval. Whether deployed using lower cost enterprise-class blade servers or rack-mounted "pizza boxes", cluster systems will deliver excellent performance and power efficiency. Unfortunately, cluster management costs are high, with a separate copy of the operating environment on each platform, and the need to replicate applications and content. This does not even address the requirement for a separate cluster file system or the limited memory available in any two-socket server.

[1] Reduced Instruction Set Computing.

Applications, like database management systems, with a large memory requirement need a scale-up SMP architecture to achieve the kinds of linear scalability and economy of scale that the data center IT budget demands. That is exactly what vSMP Foundation from ScaleMP provides, regardless of the brand name on the physical platform.

vSMP Foundation

ScaleMP's *vSMP Foundation* seamlessly aggregates multiple, low-cost dual-socket blade or rack-mounted, open systems, x86 servers into a single virtual system with linear performance scalability, superior server density, and low power consumption, in support of a variety of *Linux* operating environments. vSMP Foundation enables mid-sized enterprises and workgroups to deploy a high-performance capability outside of the scope of a traditional data center, running both memory-restricted applications and parallel workloads on a single virtual system aggregating from 4 to 32 processors, totaling up to 128 processor cores with a compute capability up to 1.75 TFLOP, and up to 4TB of shared memory, with a single point of management.

vSMP Foundation provides significant price/performance advantages for the enterprise with HPC requirements, versus an expensive proprietary, legacy SMP platform with a shared memory capability. It delivers a unique method of deploying low-cost open, x86 servers to satisfy the HPC environment and support both parallel and large memory workloads, with lower power consumption, saving valuable data center floor space by aggregating the memory bandwidth of multiple platforms and consolidating the x86 core processing capability of those platforms through the use of a high-performance InfiniBand network. This enables HPC applications running within a single VM, provided by vSMP Foundation, to exhibit near linear processing and bandwidth scalability for applications capable of taking advantage of a multi-threaded environment.

vSMP Foundation has advantages for solutions that normally operate in a scale-out architecture. By implementing vSMP Foundation, the data center can lower the deployment costs to the enterprise and simplify day-to day operations. With only one (virtual) system to manage, the IT staff will find that ease of use is increased, compared to the complexities involved in managing a cluster. There is no need for a cluster file system or the issues involved in interconnecting cluster nodes. vSMP Foundation consolidates each server's storage and networking requirements, reducing the number of drivers, HBAs, NICs, and switch ports. This reduces the number of peripheral devices that the data center has to purchase, deploy, manage, and service, increasing total system reliability, availability, and serviceability (RAS). With only one operating environment, there is no need to worry about application provisioning and installation. There is certainly no need for patching multiple operating systems. This reduces system complexity and increases system utilization, both factors contributing to an increase in the bottom line. Furthermore, with vSMP Foundation system, the data center will save significantly in reduced licensing fees.

Conclusion

If your data center is looking for an open virtualization solution to address high performance computing problems, then vSMP Foundation may be the consolidation tool that the IT staff has been looking for to improve performance through aggregating the compute, memory, and memory bandwidth resources of multiple x86 servers into a single scale-up system. It can relieve the data center of complex clustering issues, simplifying the deployment and administration of applications that require a large shared memory footprint and the aggregation of multiple CPUs. Typical environments that can benefit from vSMP Foundation are manufacturing, life sciences, financial simulations and energy.

If your enterprise is looking for better performance with higher processor utilization than you are achieving today, vSMP Foundation may be the solution that you need to reduce acquisition and operating costs as compared to proprietary systems; in fact, it will reduce your entire TCO. With a scalable architecture, vSMP Foundation can protect and extend your investments in commodity platforms.

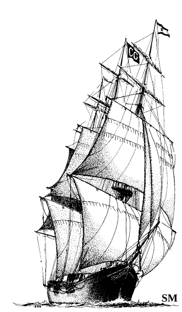

About The Clipper Group, Inc.

The Clipper Group, Inc., is an independent consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 25 years of real-world experience. A team of staff consultants augments our capabilities, with significant experience across a broad spectrum of applications and environments.

- ***The Clipper Group can be reached at 781-235-0085 and found on the web at www.clipper.com.***

About the Author

***David Reine* is Director, Enterprise Systems for The Clipper Group.** Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. He joined The Clipper Group after three decades in server and storage product marketing and program management for Groupe Bull, Zenith Data Systems, and Honeywell Information Systems. Mr. Reine earned a Bachelor of Arts degree from Tufts University, and an MBA from Northeastern University.

- ***Reach David Reine via e-mail at dave.reine@clipper.com or at 781-235-0085 Ext. 123. (Please dial "123" when you hear the automated attendant.)***

Regarding Trademarks and Service Marks

The Clipper Group Navigator, The Clipper Group Explorer, The Clipper Group Observer, The Clipper Group *Captain's Log*, The Clipper Group Voyager, **Clipper Notes**, and "*clipper.com*" are trademarks of The Clipper Group, Inc., and the clipper ship drawings, "*Navigating Information Technology Horizons*", and "*teraproductivity*" are service marks of The Clipper Group, Inc. The Clipper Group, Inc., reserves all rights regarding its trademarks and service marks. All other trademarks, etc., belong to their respective owners.

Disclosure

Officers and/or employees of The Clipper Group may own as individuals, directly or indirectly, shares in one or more companies discussed in this bulletin. Company policy prohibits any officer or employee from holding more than one percent of the outstanding shares of any company covered by The Clipper Group. The Clipper Group, Inc., has no such equity holdings.

Regarding the Information in this Issue

The Clipper Group believes the information included in this report to be accurate. Data has been received from a variety of sources, which we believe to be reliable, including manufacturers, distributors, or users of the products discussed herein. The Clipper Group, Inc., cannot be held responsible for any consequential damages resulting from the application of information or opinions contained in this report.