



HPC Computing in an Integrated Environment — Dell Delivers Virtualization with AMD

Analyst: David Reine

Management Summary

The typical consumer is always faced with an indefinite set of decisions, no matter what the product: *cost vs. functionality vs. performance vs. reliability vs. energy efficiency vs. upgradability vs. ease-of-use*, along with any other number of factors, even color. Some product decisions are easy, only involving two or three factors, others are more complicated. Take, for example, the purchase of a new car. What are the most important factors for you? No matter how much we protest, cost is always a factor. However, are we dealing with just acquisition cost, or are we more concerned with operational cost? The ever-increasing cost of gasoline is a major factor for anyone driving upwards of 20,000 miles per year. An increase of \$1 per gallon amounts to an increase in operational costs of \$20,000 per year. This is non-trivial. Do you try to keep acquisition costs down with a standard internal combustion engine or do you bite-the-bullet up front and buy a car with an electric or hybrid power plant? How important is acceleration? Do you need to go from 0-to-60 in less than five seconds? How many passengers do you need to transport? Will a two-seater sports car suffice or do you need a minivan? Will one car provide you with all of the functionality that you need, or do you need to acquire multiple vehicles to meet your driving needs? What is the cost of having a \$25,000 car sitting idle in your driveway?

The CIO of every major data center is faced with making similar decisions every day. How many servers does the enterprise need to support a unified application environment when there are transaction processing requirements, data warehouse queries, high performance computing needs, and more, facing the IT staff on a daily basis? Do you acquire x86 servers for web facing applications, while you employ proprietary, highly-scalable UNIX platforms in support of your mission-critical applications for hundreds, or even thousands of users, where reliability is of the utmost concern, or can you effectively deploy these applications on open systems x86 servers? Even if you have decided to go with an open systems deployment, what CPU is best for your solution set? Do you acquire platforms based on Intel's *Xeon* microprocessor or do you select a platform based upon AMD's *Opteron* family? Or it may come down to simply fitting your environment into whatever platform your vendor offers.

If your vendor is Dell, that is not a concern. Dell offers a multiplicity of *PowerEdge (PE)* platforms using both Intel and AMD CPUs, supporting both *Windows* and *Linux* operating environments. With Intel, Dell offers servers with the latest *Xeon 5600* (6-core *Westmere*) and *Xeon 6500/7500* (8-core *Nehalem EX*) processors, while they also offer the latest AMD *Opteron 6100 Magny-Cours* CPU, designed with a 12-core architecture to provide a maximum number of processing threads to handle the most complicated OLTP or HPC environments. In fact, Dell offers both processors within the same chassis architecture, with the *PE R810* supporting an Intel environment, while the *PE R815* supports the AMD CPU. Is AMD right for your HPC environment? To find out, please read on.

IN THIS ISSUE

➤ HPC Data Center Requirements	2
➤ The AMD Magny-Cours Architecture	2
➤ The Dell PowerEdge R815 Rack Server.....	3
➤ Conclusion	4

HPC Data Center Requirements

High Performance Computing, or HPC, is the term commonly associated with computing deployed for scientific or computational research, often on UNIX-based proprietary servers. Typically, the mathematical applications of HPC require a lot of processing power and a lot of memory, something in abundance with today's x86 microprocessors. HPC does not demand as much I/O as, say, transaction processing applications, and today HPC applications usually are run on cluster-based, large-scale, open-systems servers. Of late, some commercial applications are being consolidated with HPC applications onto open systems platforms, especially for data warehousing and transaction processing, where a higher level of I/O is required, because of the increased I/O capability of today's x86 CPUs. In general, HPC refers to the practice of aggregating computational capability in a manner to deliver significantly higher performance than a data center can achieve from a single, standard x86 platform, in order to solve large problems in business, as well as science.

What are the keys to creating an effective HPC environment? First of all, the HPC data center has a requirement for a highly-efficient virtualized architecture. It must have a very dense compute capability, with a scalable memory architecture in support of memory-intensive environments. Because of the intense power requirements of a multi-CPU compute environment, energy consumption has become a critical factor. In order to minimize the number of processor sockets, and energy, required to deliver the necessary compute density, core count has become a critical component in HPC server selection. Additionally, due to the high cost of administrative personnel, the HPC environment must be easy to use and easy to manage, in order to avoid the necessity for more expensive, highly trained administrative personnel.

HPC servers have all of the elements you'd find on any desktop computer: processors, memory, disk, operating system, applications – but more of them. High performance computers of interest to small and medium-sized businesses today are really clusters of computers. Each individual computer in a commonly configured small cluster has from one-to-four processors, and today's CPUs typically have from two-to-four cores, although, as we have seen recently, significantly higher core and thread counts are available. HPC administrators often refer to the indiv-

idual computers in a cluster as *nodes*. A cluster of interest to a small business could have as few as four nodes, with a total of 16 cores. A common cluster size in many businesses might contain 16 to 64 nodes, or from 64 to 256 cores.

The reason to have a high performance computer is so that all of the nodes can work together to solve a problem larger than any one easily can solve. And, just like people, the nodes need to be able to talk to one another to work meaningfully together. Of course, computers talk to each other over networks, and there are a variety of computer network (or interconnect) options available for the business cluster.

Your HPC cluster won't run without an operating system that can utilize all of these nodes efficiently. Two of the most popular choices for HPC are *Linux* and *Windows* (both in a variety of flavors). Linux currently dominates HPC installations, but this is due to its legacy in supercomputing, with large scale-up servers operating under *UNIX*. Your choice of O/S will always be driven by the kinds of applications you need to run on high performance applications.

A high-performance microprocessor is obviously a key consideration in the acquisition and deployment of a high performance computing solution for your enterprise, however, in an age of consolidation and virtualization, that system must be able to support commercial applications as well. Currently, the densest open systems processor chip comes from AMD, in the form of the *Opteron Magny-Cours* CPU.

The AMD Magny-Cours Architecture

In an effort to compete with the latest announcements from Intel for their *Xeon* family, AMD has introduced the newest member of their Opteron family, the Magny-Cours, with 12 cores per processor in order to improve performance for scalable applications, the processor version of an American muscle car. It also provides energy-efficiency, consuming as few as 80 Watts at 2.2 GHz, along with features to enhance management of Linux and Windows environments. The cores in Opteron are "real" physical cores not virtual cores, or hyper-threading. AMD has taken the position that physical cores provide better performance, which is especially true in CPU-intensive HPC environments. AMD provides an extensive set of on-board functionality in order to improve that performance.

- **AMD Virtualization (AMD-V) 2.0**, which offloads virtualization functionality from a

hypervisor, such as *VMware*, and onto the CPU. This provides improved I/O performance in a virtual machine (VM) by supporting I/O-level virtualization and allows direct control of devices by a VM. Opteron also supports virtualization pools and *Seamless Live Migration*, through the *AMD Extended Migration* capability, which requires all VMs to be on the same microprocessor architecture, i.e. AMD to AMD or Intel to Intel.

- **AMD-P technology** – a suite of advanced features to reduce energy use including:
 - **Dual Dynamic Power Management** to reduce idle processor power consumption;
 - **CIE Power State** turns off memory controllers and **HyperTransport 3** technology links during system idle time;
 - **AMD PowerNow! technology** to manage power utilization across processor cores dynamically; and
 - **AMD CoolCore technology**¹ to enable reduced power consumption with individual cores.

With twice the performance of AMD's previous Opteron processor, *Istanbul*, Magny-Cours supports four channels of DDR3 memory delivering up to 2.5 times the memory bandwidth (over *Istanbul*), surpassing the three memory channels in Intel's *Xeon 5600*. AMD's *Direct Connect Architecture (DCA) 2.0* supports 12 DDR-1333 DIMMs per processor to lower memory cost and better balance, key in the deployment of memory-intensive workloads including Linux virtualization, databases, and HPC, surpassing the up-to-nine DIMM slots per socket in the *Xeon 5600*. Quite simply, more memory equates to more VMs and better performance. DCA 2.0 increases the number of *HyperTransport* links between CPUs from three to four (over *Istanbul*), for faster CPU intercommunication. Magny-Cours also introduces a new *Advanced Platform Management Link* to provide an interface for processor and system management monitoring and to control system resources, including platform power consumption and cooling.

However, no data center buys simply the CPU. They acquire and deploy systems consisting of processors, operating environments, appli-

cations, and ease-of-use management tools. One company that has exhibited a determination to provide their customers with an option of the latest technologies available is Dell. In fact, Dell has introduced two rack servers based upon the same chassis, the *PowerEdge (PE) R810* based upon Intel's Xeon architecture, and a second server, the *PE R815*, based upon AMD's Magny-Cours processor, which has been designed to support the most demanding HPC environment.

The Dell PowerEdge R815 Rack Server

The PowerEdge R815 is an AMD Opteron processor-based 4-socket, 2U rack server designed for customers who need a combination of high-performance, ease-of-use, and value in HPC, virtualization, database, email, and other highly-threaded application environments. It is ideally suited as the *consolidation* vehicle for the data center, eliminating the need to waste resources on both high-performance and transactional applications. It has been designed to handle a variety of demanding workloads with outstanding speed, stability, efficiency and bottom-line value.

The PE R815 provides an excellent balance of CPU density and I/O performance, redundancy and value in a space-efficient form factor. It provides the data center with the ability to scale up to 48 processors cores using an AMD *Opteron 6100* series architecture. It also supports an advanced systems management capability with Dell's *Lifecycle Controller*.

The PowerEdge R815 is configurable with up to 512GB of ECC memory over 32 DIMM slots, at a speed of up to 1333MHz, in support of the densest HPC and virtualization environments. It has six drive bays to support hot-swap SAS, SATA and SSD drives, for a total of 3TB of internal storage. It has six PCIe G2 slots and one x4 storage slot, with five of the slots x8 and one x4. It has two integrated dual-port Broadcom 5709C 1Gb NICs and comes with a variety of optional network controllers. The PE R815 also has an interactive LCD screen to facilitate administration with Dell's *OpenManage* software, improving availability. It also supports multiple operating environments, including a wide variety of Microsoft *Windows Server 2008* implementations and Microsoft *SQL Server 2008 R2*. In addition, it supports *Red Hat Enterprise Linux*, *Novell SUSE Linux Enterprise Server*, *Citrix XenServer*, *VMware ESX 4.1*, and *Sun Solaris*.

¹ AMD CoolSpeed technology reduces p-states when a temperature limit is reached to allow a server to operate if the processor's thermal environment exceeds safe operational limits.

The PE R815 reduces power consumption via two hot-pluggable redundant power supplies. Your application requirements will determine if you require redundancy at the power supply.

The PE R815 supports multiple users, and higher transaction levels. With a higher level of processor utilization from expanded virtualization capabilities, the R815 provides the data center with an ideal platform for consolidation. It enables the IT staff to reduce server count, improve utilization of computational and memory resources, and enables the data center to reduce the number of software images being maintained. With the Opteron 6100, Dell can help the staff reduce memory costs with more DIMM slots using less expensive, smaller DIMMs.

The PowerEdge R815 also includes power management features such as voltage regulators, power regulating processors and an interactive LCD screen for easy access to power consumption information. Along with the server's internal design, these features can simplify and aid in the management of power within your IT environment.

That design includes low-flow fans that spin in accordance to server workload demands, helping to reduce unnecessary noise whenever possible. In addition, the shrouding and layout of the internal components help to keep the server cool.

Reliability

The PE R815 is designed for reliability and ease of use, with features ranging from robust, metal-hard drive carriers and embedded diagnostics to an interactive LCD display and Industrial-quality materials. The PE R815 also includes Dell's Lifecycle Controller with dual internal Secure Digital (SD) modules to provide failover at the hypervisor level.

Dell has established a simple reliability goal: Deliver quality products that stand the test of time. To help meet this goal, Dell is improving reliability processes by:

- Employing robust validation and testing processes for durable product design;
- Verifying that each supplier meets Dell's stringent quality standards;
- Utilizing robust, durable industrial materials to enable long product lifecycles;
- Introducing a *Unified Server Configurator (USC)*, which helps to minimize downtime by offering embedded and persistent diagnostics with no media required; and
- Ensuring that every fully configured server is

tested before it leaves the factory.

Management

Along with its adherence to open standards, Dell's systems-management portfolio focuses on simplicity, efficiency, cost containment, and cost reduction. Additionally, Dell has integrated its management solutions with third-party offerings from BMC, Microsoft, and VMware, while providing connections to HP *Openview* and IBM *Tivoli*, as well as solutions from Symantec and CommVault, all aimed at providing the enterprise with a holistic focus across the solutions stack.

Integrated on the server, the Lifecycle Controller helps to simplify management tasks by performing a comprehensive set of provisioning functions such as system deployment, system updates, hardware configuration, and diagnostics in a pre-OS environment using the USC. This can help eliminate the need to use and maintain multiple pieces of disparate media.

The PowerEdge R815's interactive LCD display is positioned conveniently on the front of the server for easy setup, monitoring, and maintenance. Plain-language diagnosis and a programmable messaging system can help users address issues quickly, while a color-based alert system (blue for normal operation; amber for a problem) simplifies day-to-day monitoring.

Dell's *OpenManage* systems-management portfolio also includes the *Dell Management Console*, which comes with every Dell server, and provides IT managers with a consolidated view of their entire IT infrastructure.

Conclusion

Along with rising energy costs, limitations on space, power and cooling capacity present enormous challenges for enterprise IT. The PE R815 server drives energy efficiency as a design standard while delivering the consolidation and virtualization performance you need to meet cost and environmental goals. Dell's energy-efficient system design, built with *Energy Smart* technologies, can help to manage power in your HPC environment helping to preserve the life expectancy of your data center. If you are looking for a HPC platform to help manage your consolidation and virtualization, take a look at the Dell PE R815.



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 a Senior Contributing Analyst for The Clipper Group. Mr. Reine specializes in enterprise servers, storage, and software, strategic business solutions, and trends in open systems architectures. In 2002, 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.

Disclosures

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.

After publication of a bulletin on *clipper.com*, The Clipper Group offers all vendors and users the opportunity to license its publications for a fee, since linking to Clipper's web pages, posting of Clipper documents on other's websites, and printing of hard-copy reprints is not allowed without payment of related fee(s). Less than half of our publications are licensed in this way. In addition, analysts regularly receive briefings from many vendors. Occasionally, Clipper analysts' travel and/or lodging expenses and/or conference fees have been subsidized by a vendor, in order to participate in briefings. The Clipper Group does not charge any professional fees to participate in these information-gathering events. In addition, some vendors sometime provide binders, USB drives containing presentations, and other conference-related paraphernalia to Clipper's analysts.

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.