THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2011038 December 20, 2011

HP Raises the Bar for Hyperscale Computing While Lowering Energy Requirements and TCO

Analyst: David Reine

Management Summary

Choice. Whenever we go shopping, we crave it, so that we can purchase an item, big or small, that suits our needs. We have seen any number of "Big Box" stores that cater to larger families with super-sized boxes of cereal and jugs of ketchup. We see multiple brands for each product, so that we can choose whichever one best suits our taste. We also see a significant amount of choice in the cars that we purchase. For some of us, that sleek sports car with 0-to-60 MPH acceleration in less than six seconds fits our desire; for others the decision will be made based upon the number of passengers (including the family dog) that need to be transported. Still others will be concerned about fuel economy, so the intelligent hybrid or all-electric vehicle might be the most suitable acquisition.

Determining what is best for a specific situation is nothing new for the CIO or data center decision maker. The enterprise data center is being overrun by a virtual plethora of servers, deployed to execute any number of applications. Some of these applications require a significant amount of processing power to accomplish their tasks while others require more memory in order to support the number of virtual machines (VMs) needed to satisfy an environment that recently has been consolidated and virtualized. Still other applications are dependent heavily on I/O throughput, in order to satisfy user needs. In most enterprise data centers, we have seen the deployment of hundreds (even, thousands) of open systems servers based upon the x86 architecture, either Intel *Xeon* processor or AMD *Opteron* CPU. Both use a CISC (Complex Instruction Set Computing) architecture to execute a myriad of compute instructions. Fortunately, not all applications need to execute all of these commands. For these applications, a RISC, or Reduced Instruction Set Computing, environment is more appropriate. Unfortunately, most of these CISC and RISC processors consume a lot of energy, upwards of 100 watts for just the CPU. Many even consume up to 80 watts in idle mode. This may not sound significant for a small data center with less than 10 servers deployed; however, it becomes very significant for the enterprise data center with thousands of multi-socketed servers humming around the clock. Because of a lack of choice, many CIOs are paying for instructions and energy that provide no real return (benefit) for these "extras". So, where can we find a CPU architecture that will provide the performance they need, without the extreme energy consumption? Try looking in your pocket!

The processors used extensively in many consumer electronics are called *ARM*, for *Advanced RISC Machine*. They can be found in mobile phones, digital media, and music players, and hand-held game consoles, devices that do not require a lot of horsepower and consume energy in sips instead of gulps. If an ARM CPU could be adapted for a data center server environment, that might be very good. In fact, that is exactly what Hewlett Packard (HP) has done with its *Redstone* server, with the help of Calxeda's ARM system-on-a-chip. To learn more about the latest in server technology, please read on.

IN THIS ISSUE

Enterprise Data Center Requirements

Today's enterprise data center faces several critical problems. Because of the server sprawl that has overwhelmed the entire enterprise, not just the data center, the CIO is facing the daunting task of reducing complexity, lowering energy consumption, and reducing the total cost of ownership (TCO) of IT infrastructure. In an attempt to lower the complexity of that infrastructure, many enterprises have deployed a strategy of consolidation and virtualization in order to make better utilization of its IT resources and floor space. This strategy enables them to share many of the siloed resources that were going unused, including processing power, storage, networking, and management, and help them lower the TCO of the data center.

Power consumption is another factor that has to be reduced, but usually not at the expense of performance. Each new generation of open systems servers, based upon the x86 architecture, improves processing power and reduces energy consumption. Nonetheless, we are now living in a new, connected era and we rapidly are running out of both floor space and energy in the data center as we deploy hundreds, even thousands, of new servers to process a myriad of new digital content every day. This content might be made up of millions of photos, emails, texts, tweets, and videos daily. In fact, HP has projected that, by 2015, more than *98 billion* mobile applications will be downloaded to more than *seven billion* mobile devices worldwide. Consolidating servers has proven effective, on an evolutionary scale, in conserving energy. However, if the enterprise intends to continue growth on the current path, a revolutionary change to the server paradigm is essential. We simply cannot continue to throw money and more power at this burgeoning problem. A lower-energy-consuming processor is required if we expect to continue our thirst for information that requires continuing server and storage growth to drive our highly-connected business environment. Cutting back one server at a time may be approaching from the wrong end of equation. Evolution will consume us. We need revolution at large scale. We are facing a ten times growth in server count, from a mere 50 *million*, to a ridiculous half a *million* over the next three years as we connect everyone to everything. *Can you now imagine a doubling of the energy requirement or, perhaps, as much as a ten times increase, depending on your assumptions?*

We need to improve the power usage effectiveness (PUE) of energy in our enterprise data centers. PUE is a measure of how efficiently an enterprise uses its power; specifically, how much of that power is actually being used by the IT infrastructure (in contrast to cooling and other overhead). **We need to make a significant leap in energy efficiency, if we expect to meet the growth demanded for servers and storage over the next decade.**

All Workloads Are Not The Same

Some enterprise applications do not require the full set of instructions that are available in an x86 open systems architecture. These include applications with light scale-out workloads, simple content delivery (i.e. static web servers with relatively little processing for each of thousands (or millions) of requests, and video servers fetching and shipping unmodified video). Other examples include big-data-oriented simple search systems, simple database searches on in-memory databases, along with other applications with lightweight CPU requirements. In general, lightweight applications (i.e., not processing intensive) may not need the complex instructions set and neither will most open applications (i.e., not "Wintel"[1]).

What are left are those applications that aren't running on Windows (and, typically, are running on Linux), especially those that might be classified as High-Performance Computing (HPC), whether run in-house or offered by Internet service providers. HPC applications tend to run in pieces in parallel (i.e., on many cores) and the "results" often are assembled from the "answers" that come back from processing each "piece" of the problem being solved. These tend to be mathematical computations and analyses that scale well, i.e., they can run on many cores at the same time, each one operating independently. For many businesses, these applications tend to run in *hyperscale* (very large) proportions.

One such hyperscale arena is in the financial markets, where the volume of data continues to rise exponentially, and traditional scale-up and scale-out models struggle to keep up with demand without having to increase cost and power

[1] Windows on x86-based servers. In general, if an application requires the CPU native instruction set, either because of operating system, compiler, or other code written especially for (i.e., only for) that CPU, then the discussion that follows does not apply.

consumption significantly. To meet the computing needs of these environments, and others, new, low-energy consuming systems need to be considered, especially if the data center is challenged to live within existing power supply limitations. Extreme low-energy servers are also attractive for cloud computing and social media environments. Energy-conscious social media companies, such as Google and Facebook, where IT *is* the business, are expected to take advantage of this new technology. To them, total cost of application (and/or data) delivery is the primary TCO measurement.

New thinking and industry leadership are required to change from the traditional IT paradigm (i.e., x86 and other CISC servers), whenever it is economically advantageous. For many hyperscale applications, performance may not be the primary concern; electric power is. The work *can* be done with traditional processors, but more energy (and also floor space and potentially more CapEx cost) are required. **If you have mathematically-centered applications that are reaching hyperscale proportions, you should be considering alternative architectures.**[2]

TCO issues, including space and energy, must be addressed to unlock the savings potential within the data center. ***What difference does it make in the cost/performance of your computing architecture, if you do not have the power to turn them on?*** In order to meet the needs of many large enterprises, HP, which proclaims itself as the leader in scale-out computing systems, has teamed up with Calxeda to deliver this new data center paradigm of low-energy servers.

Who (or What) is Calxeda?

Calxeda is a startup server company that is the innovator behind the world's first five-watt server. It is a semiconductor company that does not believe that bigger is better, at least when it comes to hyperscale building blocks. Rather, Calxeda believes in using the *right tool* for the right job. In this case, that tool is an *EnergyCore ARM Cortex* processor[3], designed first for power efficiency. Calxeda's processor delivery system can scale up to thousands of server nodes. The ARM Cortex RISC CPU is dependable, reliable, and responsive enough for use in a mission-critical enterprise data center, although not for CISC-specific application environments, as discussed previously. Instead, Calxeda's processor delivery system is focused on doing offline analytics, web applications, middle-tier applications, and file/storage serving, with a keen emphasis on performance per watt.

Despite Calxeda's paradigm-shifting innovation, it realized that enterprise data centers weren't going to bet its future on a startup. Thus, to propagate its architecture; it needed to partner with an established player, a leader in data center IT infrastructure. In order to take the risk out of an innovative RISC architecture, Calxeda turned to HP, which then selected Calxeda as its first partner in its *Project Moonshot*.

What is Project Moonshot?

Project Moonshot is an agnostic architecture that provides federated management, fabric, storage, networking, and power/cooling, in a rack-oriented infrastructure, to simplify the scaling of hyperscale application delivery. It is comprised of three distinct elements:

1. The HP Redstone Server Development Platform (discussed on the next page),
2. HP's new *Discovery Lab*, and
3. HP's new *Pathfinder Program* for partner collaboration.

Project Moonshot has its roots in HP's *ProLiant* servers (x86), which use Intel's Xeon and AMD's Opteron processors. However, unlike an x86 architecture that is capable of being configured with tens of servers per rack, Project Moonshot is an environment that enables thousands of cores to be deployed in a single rack, with a processor-neutral architecture and a federated infrastructure that scales seamlessly as additional servers are added, all while sharing resources such as power/cooling and storage. However, hyperscale number-crunching deployments tend not to use virtualized cores, as each parallel piece of the application tends to consume the entire core. HP has partnered with several other companies to break new ground for hyperscale computing, changing the IT paradigm with regard to environments such as cloud services and on-demand computing. As discussed earlier, many enterprises with hyperscale requirements rapidly are approaching a "tipping point" in energy consumption, one that will force fundamental changes in systems architecture.

[2] Conversely, if your applications are not processor intensive, they would not be well suited to this new architecture.

[3] A single Cortex processor represents one server.

Exhibit 1 – A View Inside One of HP's Redstone Rackmount SL6500 Servers

Source: HP

Project Moonshot is a multi-year program that takes advantage of HP's extensive experience in deploying enterprise-level private and public clouds and also recent experience with a new category of servers, using low-energy computing resources that are normally found in mobile devices, such as your mobile phone. With HP's *Converged Infrastructure* as its foundation, Project Moonshot pools resources in a highly federated environment to reduce space and energy requirements, as well as management complexity. By pooling resources, HP can take advantage of low-energy processing to build a scale-out cluster to excel at highly parallelized, low-transaction workloads while reducing the TCO of the IT infrastructure. HP projects that solutions build on Project Moonshot will consume up to 89% less energy and occupy 94% less space, thus lowering TCO by 63%. The HP *Redstone Server Development Platform* (discussed next) is said to be the first in a line of extreme low-energy server platforms from HP.

HP Redstone Server Development Platform

Redstone is an offshoot of HP's leading low-energy innovations, such as the HP *Data Center Smart Grid*, and has been deployed to further the development of low-energy server technology. Initially, Redstone has been configured with an array of *EnergyCore* ARM Cortex processors from Calxeda, delivered as multi-processor boards that fit into HP's innovative delivery solution, lowering the power and space requirements for scale-out computing. While Calxeda makes "low-energy processor delivery systems", HP delivers "hyperscale processor solutions". In addition to the processor delivery systems (boards) from Calxeda, the EnergyCore CPU includes an *EnergyCore Fabric Switch*, an *EnergyCore Management Engine*, providing autonomic power management and fabric routing, along with I/O controllers. Each CPU runs at up to 1.4GHz and supports floating-point operations.

The Calxeda CPU (node) is robust, with four interconnected ARM cores to remove network bottlenecks, integrated with 4MB of L2 cache and integrated, high bandwidth memory controllers. A single CPU is extremely efficient, consuming only five watts of power when running at full speed, which drops to less than .5 watts when idle. Each processor has integrated systems management and a high-performance fabric to converge inter-node communication, I/O, storage, and networking. The EnergyCore ARM Cortex CPU is easy to use with standard drivers and standard interfaces.

The Redstone Server Development Platform decouples scale from complexity and consumes only a fraction of the power required by conventional x86 processors. Because of this low-power requirement, Redstone can achieve unparalleled density through an ultra-small form factor, enabling a design to support in excess of 2,800 CPUs in a single rack, reducing server sprawl, cabling requirements, switching, and peripheral devices. These processors are configured within HP's *Scalable System SL 6500s* chassis. A single 4U SL6500 supports up to four 1U-equivalent trays[4], each configured with 18 quad-node compute cartridges totaling 72 Calxeda EnergyCore ARM servers per tray, with 288 processors per 4U chassis. (See Exhibit 1, at the top of the previous page.) Each CPU supports up to 4GB of ECC memory, with integrated management, and can be diskless or support up to four SATA drives. The SL6500's scalable enclosure shares pooled power via four common slot power supplies, shared cooling with eight shared fans in an *N+1* configuration. It also has up to 192 SSDs or 96 2.5" SFF HDDs and an integrated, configurable network fabric, with up to sixteen 10Gb uplinks.

As a result, Redstone reduces floor space, energy consumption, and software licensing, and thus lowers complexity by 97%, according to HP. The initial Redstone platform will be available to customers in the first half of 2012, with follow-on products to be delivered, as possible alternative platforms.

HP Discovery Lab

In support of the Redstone platform, HP will open the HP Discovery Lab in Houston in January, to help potential clients navigate the journey to determine application suitability in an extreme low-energy environment. HP's prospects and customers can access the lab to experiment, test, and benchmark applications, to determine which applications will unlock the potential of extreme low-energy servers. HP also will open additional sites in Europe and Asia, later. All of these labs will enable both remote and onsite access, so that prospects and customers can work directly with HP engineers and industry peers, in open collaboration with other program participants through a secure, social media portal.

HP Pathfinder Program

In order to be able to deliver a complete ecosystem for Project Moonshot, HP has developed a partner program called *Pathfinder* to establish industry standards and best practices. This program covers silicon development, systems, operating environments, and ISVs, among others, to solidify efforts in improving power efficiency in the enterprise data center.

Conclusion

First, HP is more than just another purveyor of clone PCs and servers to a thirsty data center. It is an innovative leader in enterprise data center architecture, well positioned to take the enterprise on a journey to energy efficiency. Second, Project Moonshot is not just another server with an ARM chip instead of an x86 processor; it is a complete ecosystem integrated to reduce data center complexity and cost.

HP's Redstone platform truly represents a paradigm shift in IT infrastructure. However, it merely is the first step in a long journey. It just may well represent a niche market today (for the HPC crowd), but it may be the savior of the enterprise data center in the years to come. Stay tuned, as we follow HP on this journey.

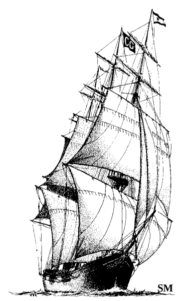

[4] These are half-width 2U trays; that's how each is equivalent to 1U.

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.