



With Enhanced POWER7 Systems, IBM Raises the Ceiling, Again

Analyst: David Reine

Management Summary

It is clear that we live in a land of the *survival of the fittest*. We see this everyday in sports as well as our commercial endeavors. We have just witnessed the NCAA basketball playoffs where 68 teams engaged in hand-to-hand combat until there was only one team left standing. We can drive down the street to the nearest strip mall and view the rows of empty stores, where thriving businesses had once stood. Where once we had our choice of retail electronics stores, such as Circuit City or CompUSA, we now often find empty buildings. The same is true for retail stores where we have seen several chains, such as Border's and Blockbuster's, closing many of their outlets or entirely. Whether it is the recession or an inability to compete, the consumer has fewer options. Even when you decide to go clothes shopping, you find stores, such as Ann Taylor and Abercrombie & Fitch, disappearing in front of our eyes. Thank goodness for Internet shopping. You also need to call ahead when going out to eat to be sure the restaurant is still open, with so many long-standing disappearing over the last few years. One might say that only the fittest survive, but that doesn't help you when your favorite retailer or restaurant is no more.

This same scenario is certainly not news to the CIO of any enterprise data center, especially in the data center with a mission-critical infrastructure based upon the UNIX operating system. Where once he had a lineup of salesmen knocking on his door, today there are few. Over the past 40 years, the data center has seen a variety of UNIX systems come and go, as well as a variety of vendors who can no longer be found. At one time, we had *System V* from Bell Labs, *BSD* from Berkeley, *Ultrix* from Digital Equipment, *IRIX* from Silicon Graphics, *AIX* from IBM, *Solaris* from Sun, and *HP-UX* from HP, as well as many others. Today, in fact, there are really only three systems vendors peddling UNIX on specialized architectures: AIX on *POWER7* on the IBM *Power Systems*, HP-UX on Intel *Itanium* on *Integrity*, and Solaris on *SPARC* for the *UltraSPARC* family, once from Sun, now from Oracle. However, what does the future hold for these three?

In fact, the futures of HP-UX on Itanium and Solaris on SPARC have also been questioned. Oracle has terminated future software efforts on Itanium and Intel continues to put mission-critical RAS features into its primary enterprise processor, *Xeon*, while Oracle's commitment to the SPARC architecture also has been doubted. How many resources are they willing to invest? No one, however, is questioning IBM's commitment to AIX and *POWER7* components to their *Power Systems* family. IBM continues to enhance the consolidation and virtualization features of *POWER7*, as well as improving the RAS characteristics, in order to deploy mission-critical, workload-optimized systems. With its near 50% revenue market share and continued share gain, IBM claims that it continues to lead in the \$16B UNIX market. *POWER7* set the standard when it was announced last year. Now, IBM has taken their leadership position another giant step forward, not only in hardware, but in software as well. In order to learn more about IBM's 2011 enhancements to the Power family, please read on.

IN THIS ISSUE

- **Mission-Critical Data Center Needs 2**
- **Today's UNIX Market 2**
- **IBM POWER7 Family 3**
- **Conclusion 5**

Mission-Critical Data Center Needs

The CIO of any modern, mission-critical, enterprise data center faces a raft of problems, starting with the sprawl of under-utilized servers that proliferate not only throughout the data center, but also throughout the enterprise as a whole. This sprawl contributes to an increase in the total cost of ownership (TCO) of the IT infrastructure, not only due to the inefficient use of the micro-processors used in that infrastructure, but also due to the energy wasted in operating the server platforms and cooling the IT environment. The data center also is spending too many Dollars (or Euros) for the administration and maintenance of this heterogeneous mix of under-utilized systems.

The data center staff is looking to optimize the workload and increase the performance throughout the enterprise, using higher levels of consolidation and virtualization in order to reduce the number of servers that it must manage, and to simplify the management and administration of the infrastructure. Calculation-intensive applications can always take advantage of all of the processing power that the server possesses.

Make no mistake about it – the lifecycle management of the data center consumes an inordinate amount of IT staff resources. The lifecycle management of the infrastructure includes the ordering, configuration, deployment, update, and patching of the mission-critical servers that are the lifeblood of many enterprise data centers. The simplification of infrastructure management remains a key factor to the ability of the staff to lower the TCO and remain on budget. Simplification also will improve the price-performance ratio, which may be one of the metrics on which you are measured.

Frequently, the consolidation and virtualization of data center servers requires the upgrading of the existing infrastructure to newer models in order to reduce not only the number of platforms in the data center, but their footprint and energy consumption levels, as well. Improving the energy efficiency of the data center, reducing the amount of energy consumed to run and cool the data center, not only makes your enterprise look good, it makes you look good.

Keep in mind that the CIO is not looking for just a better platform, but a better solution. Coincidentally, better platforms provide the resources required to deliver a better solution. All of those solutions will benefit from binary compatibility. Binary compatibility of applications between old and new systems helps to control data center

operating costs. All things being equal, requiring a migration in order to move forward may be more time consuming and expensive than one of the more straightforward approaches, if you have an option.

In the past, those better solutions have all been based upon the UNIX environment, with all of the reliability, availability, and serviceability (RAS) features that did not exist in any of the commodity solutions that were perceived to be less expensive to procure and operate. Moreover, there were a virtual plethora of proprietary UNIX solutions to choose from. Most of those mentioned earlier now belong to the historical scrapheap of IT systems. The three UNIX solutions that have maintained credibility and survived for the past three decades are *AIX* from IBM, *HP-UX* from HP, and *Solaris* from Sun. Where do they stand today?

Today's UNIX Market

The picture of the UNIX market today is the same as the newest rage in IT: *it is in the clouds*. In fact, all things may not be equal! The HP platform is based upon Intel's *Itanium* processor, an *EPIC*¹ CPU noted for its high performance and RAS characteristics. Unfortunately, for HP, as well as other vendors incorporating Itanium, Intel has been putting many of these same RAS features into their *Xeon* family, especially prevalent in the *Nehalem* and *Westmere* CPUs. This, apparently, is enough of an indicator to convince Oracle to cease development of its new applications on the Itanium architecture, if it needed another incentive to tweak HP. HP has insisted that they will continue to support existing versions of the Oracle programs on today's and future versions of their *Integrity* platform, but is now using social media to get their Itanium customer base to pressure Oracle to reverse its decision. Because Microsoft and Red Hat have already taken a similar position, this is unlikely. The clouds are darkening for HP-UX, as a result.

Oracle, on the other hand, has problems of its own. Many customers have questioned Oracle with their wallets, as they have already moved from *SPARC* servers to HP's *Integrity* or IBM's *Power Systems* servers. In fact, some have rationalized that Oracle's decision to stop development on Itanium may be simply one more attempt to stop the bleeding. Oracle claims to be committed to *SPARC* and has released a roadmap

¹ Explicitly Parallel Instruction Computing.

indicating extended support for SPARC. Unfortunately, some continue to believe that Oracle may not be totally committed to SPARC. One indication of this may be their *Exadata Database Machine*. The Exadata storage server now has been deployed on a pair of Xeon 6-core processors instead of SPARC. Oracle's dedication to SPARC can be compared to the commitment of the chicken to your breakfast eggs and the pig to your morning bacon. The chicken is involved, but the pig is totally committed! Is Oracle simply involved in their SPARC offering, or are they totally committed?

However, no one is questioning IBM's commitment to AIX or its Power systems. They are in firm control of the UNIX market. Last year's announcement of *POWER7*² continued IBM's path of yearly innovation and was adopted readily by their base. In addition, the overall UNIX community, looking for a safe harbor for their mission-critical applications, continued to look to AIX for safe and innovative solutions, delivering cradle to grave support. With as many as 3500 mission-critical data centers migrating from SPARC and Itanium to *POWER3*, it is becoming apparent that a new market segment is emerging, one for "smarter" and "big data" applications, where IBM seems likely to do very well in their competition with companies such as Dell, HP, and Oracle.

These days, IBM is acting as a very confident vendor, unwilling to concede anything to their competitors, and that includes the high-end Xeon servers. In fact, IBM is staying on course, meeting their roadmap guidelines. They have now released a faster *POWER7* processor in the *Power 755*, and the *Power 750* servers, the same systems that constituted the nodes of the *Jeopardy*-champion, *Watson*, and new *POWER7* blades⁴.

IBM POWER7 Family

IBM is taking advantage of the improved speed and functionality of *POWER7* to enhance their *Power 750* server, *Power 755* high performance engine, and its *Power* blades. Today's *POWER7* may not be faster than a speeding bullet, but it is more powerful than each *POWER7*

CPU used in *Watson*. *Watson* used *POWER7* processors running at 3.55GHz in multiple *Power 750s*. The new *POWER7* CPU runs at 3.73GHz. Using IBM's *rperf*⁵ ratings as a guide for the performance of a processor with 32 threads – a *Watson* node is rated at 91.96, while a new *POWER7* node is now rated at 101.57, about 10.5% faster. With two processors installed (64 threads), a *POWER7* system has an *rPerf* rating of 134.11, with a rating of 251.45 for a four-processor *Power 750* system with 128 threads.

As was very evident with *Watson*, new, higher-speed *Power 750* servers can access data faster and improve response times, getting results much faster than running the same calculations on a server based upon Intel's Xeon CPU. Now, with the assistance of Nuance, IBM is combining their own *DeepQA* technology – the technology used for *Watson*'s unstructured text analytics and natural language processing – with speech recognition and clinical solutions from Nuance to adapt the innovations demonstrated by *Watson* to the healthcare industry, to assist medical professionals in the diagnosis and treatment of patients.

POWER7 systems are optimized to address the needs of many workloads, with *DB2* to take advantage of the raw performance of *POWER7* and *WebSphere* to take advantage of the multi-threading capabilities. Combined with *PowerVM* virtualization, this enables the data center to do more workloads with the same number of servers, or the same number of workloads with fewer servers. It also will enable the IT staff to save on both floor space and energy costs, helping to lower the TCO of the IT infrastructure, and, as a result, helping to reduce greenhouse gas emissions. Extensions to IBM's *System Director* platform management software last year, along with this year's introduction of the *System Director Management Console (SDMC)*, also help to lower management costs. *SDMC* is the next generation management appliance for *Power* Systems. It replaces IBM's *Hardware Management Console (HMC)* with the same form factor. It can manage all *POWER6* and *POWER7* servers and blades, utilizing *System Director*.

IBM's *PowerVM* virtualization enables the IT staff to get more workloads from the *POWER7* by improving the utilization of each processor and helping to manage the most demanding emerging applications, being deployed for healthcare management, financial services, and scientific re-

² See [The Clipper Group Navigator](#) dated March 4, 2010, entitled *Optimizing the Enterprise Datacenter – POWER7 Powers a Smarter Infrastructure*, available at <http://www.clipper.com/research/TCG2010006.pdf>

³ According to statistics from IBM.

⁴ IBM is also previewing the *Power 775*, touted to be a supercomputer node and the pinnacle of *Power* technology.

⁵ Relative Performance.

search. This includes a *POWER Hypervisor* and options for *PowerVM Standard* or *Enterprise Edition*. By consolidating workloads from multiple applications running under *AIX*, *IBM i*, and *Linux*, onto a single platform, the data center can save even more space and energy along the way. By enabling the IT staff to utilize the power of POWER7 fully, PowerVM allows the Power systems to handle unexpected workload peaks.

IBM has continued the functionality of their POWER7 systems to the new, enhanced models, including:

- *Active Memory Expansion* to enable more workloads to be completed with existing server resources;
- *Light Path Diagnostics* to readily diagnose hardware problems, to provide an intuitive means to positively identify failing components;
- *Advanced RAS features* such as IBM *Chipkill* ECC detection and correction, processor instruction retry, a service processor with fault monitoring, and hot-plug bays and power;
- *The inclusion of SSD drives* to further enhance the reliability of POWER7 servers with: longer life, no head crashes, no moving parts, thus very low failure rates and no noise; SSDs produce higher performance⁶ with better energy efficiency⁷, thus a lower cost with less heat dissipation;
- *PowerHA* for the data center's high-availability requirements;
- *IBM Systems Director Active Energy Manager combined* with POWER7 *EnergyScale* technology, to help make the IT infrastructure more energy efficient. It has *Intelligent Energy* management features to operate dynamically at a higher frequency to improve performance per watt, when environmental conditions allow, or to conserve power to operate at a reduced frequency for significant energy savings.

Enhanced Power 750 Express

The *Power 750 Express* has now been further optimized for the most challenging analytics workloads. In addition to a faster POWER7 pro-

cessor, IBM is now delivering the new *Systems Director Management Console* appliance to enable the data center to operate both Power servers and blades with a unified, intuitive interface for managing both physical and virtualized system resources. They have also qualified two new BNT Ethernet switches for top-of-rack deployment. These switches have now been integrated closely with Power servers to support optimized workloads.

The Power 750 Express is a one-to-four socket server (4U) running at speeds of up to 3.72 GHz in an eight-core configuration. It has 256 KB of L2 cache per core and 4MB of L3 cache per core. The Power 750 supports up to 512GB of RDIMM DDR3 memory and up to 8 small form factor (SFF) SSD drives or up to eight SFF SAS drives. The faster POWER7 CPUs offer more than three times the performance of comparable 32-core servers, such as Oracle's *SPARC T3-2* server, and more than twice the performance of HP's *BL890c i2*⁸, all with just eight cores. In addition to the higher performance, the Power 750 Express is *ENERGY STAR* compliant, retaining all of the functionality of the previous offering.

Enhanced Power 755 Express

The IBM Power 755 is a high-performance computing cluster node with 32 POWER7 processor cores that is optimized for running highly-parallel, computationally-intensive workloads and algorithms, such as weather and climate modeling, computational chemistry, physics and petroleum reservoir modeling.

The Power 755 server is a four-socket, 4U rack-mounted server that offers 32 POWER7 processor cores running at a frequency of up to 3.6 GHz. The data center can configure up to 64 GB of memory for each processor, up to a maximum of 256GB of memory in a four-processor system. The IBM Power 755 compute node is designed for organizations that require a scalable system with extreme parallel-processing performance and dense packaging.

The Power 755 offers enhanced signal and image processing through support of the *AltiVec* instruction set and extended *VSX SIMD* acceleration, which can execute up to eight single-precision or double-precision floating point operations per clock cycle per core to improve fine-grained parallelism and accelerate data proces-

⁶ SSDs have a similar write performance but faster access and read performance than HDDs.

⁷ SSDs consume 3W while active; 1W in idle mode.

⁸ Based upon SPECint_2006 performance of a 32-core Power 750, 32-core SPARC T3-2, and a 32-core HP Integrity BL890c i2.

sing. Using 12X *InfiniBand* adapters, up to 64 Power 755 nodes, each with 32 cores, can be clustered together, providing up to 2,048 POWER7 cores.

Enhanced POWER7 Blades

IBM's newest Power blades enable the data center with BladeCenter to consolidate further and to virtualize the IT infrastructure, in order to better utilize enterprise resources. These new blades, the *BladeCenter PS703* and the *BladeCenter PS704 Express*, help to simplify the deployment of new applications. They enable the IT staff to improve efficiency, automate processes, reduce energy consumption, and scale easily as demands on the infrastructure increase.

The PS703 is a 16-core, single-wide blade server, while the PS704 is a 32-core, double-wide blade server. The latter allows the IT staff to pack servers more densely in a smaller space, in order to minimize the number of racks and, thus, reducing data center sprawl. In fact, based upon rPerf results⁹, the PS704 delivers 60% higher performance, with twice the number of cores, as the *BladeCenter PS702*, the prior doublewide blade running at 3.0GHz, within the same space and energy envelope. At the same time, the PS703 delivers better price-performance than comparable Xeon blades, based upon a total cost of acquisition comparison to the HP BL620c prepared by IBM. Both new BladeCenter blades support high levels of consolidation to better use IT resources with more energy-efficient economics.

Conclusion

With this latest upgrade to the POWER7 architecture, IBM is redefining performance expectations for the IT infrastructure. They are enabling data centers to deliver enterprise applications and services faster than ever before, and at a lower cost. They are enabling higher rates of consolidation and virtualization, thus improving on the utilization of all IT resources: from CPU to memory to I/O to storage.

With the availability of faster processors, combined with solid-state disks, POWER7 systems can accelerate all of the enterprise transaction processing requirements as well as all web-facing applications. With POWER7 servers, IBM has delivered on their promise of Smarter Com-

puting. What may be even more significant is that IBM can be counted on to continue to deliver mission-critical systems to the enterprise data center. POWER is not going away anytime soon! AIX is not going either! If you are looking to reduce the TCO of your IT infrastructure and are concerned about your existing platforms, take a long look at IBM's POWER7 servers. It can help you retain your peace of mind, and it is here to stay.



⁹ PS704 rPerf = 251 for 32 cores; PS702 rPerf = 154 for 16 cores.

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.