THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2009040 September 21, 2009

IBM Takes Command of the UNIX Data Center — POWER7 Enables Growth, Lower Costs

Analyst: David Reine

Management Summary

Striving for excellence is a human trait shared by many, as we all try to be the best that we can in at least one area under our control. *Achieving* excellence is a little harder to accomplish; it requires an amount of hard work and dedication that only a select few are willing to deliver. *Improving* on excellence, on the other hand, requires that rare individual who sets his sights on being the best in the world at whatever he attempts and continues to work harder than everyone else, even after he has arrived at the pinnacle of his quest. Individuals like Olympic athletes Michael Phelps and Usain Bolt each set world records (in swimming and track), yet each continues to train even harder to break their own records and reap the rewards of these continuing efforts.

This same quality of continuing to improve on success is an essential requirement for every enterprise data center looking to improve upon the performance of its IT infrastructure, ensure the security and reliability of its environment, and continue to lower the total cost of ownership (TCO) of that infrastructure in the face of increasing demands. The deployment of new applications on new servers and the continuing explosion of data, which tends to be doubling every 12-to-18 months, are putting a strain on the budgets of every enterprise data center around the globe. Programs are being implemented to consolidate and virtualize both servers and storage to reduce the TCO and preserve valuable resources, both human and natural. By reducing the number of physical servers populating the data center, the CIO can reduce the number of systems administrators required to drive the IT infrastructure, as well as **reducing the amount of energy** necessary to power the data center, and **the amount of floor space** required to house it. **These last two points are especially critical as enterprise data centers approach maximum capacity in both of these categories.** In fact, if either is exceeded, the enterprise may be forced to build out a brand new data center at a cost of millions of dollars.

For decades, UNIX systems vendors have been providing continual and timely upgrades, enabling customers to run more applications, respond faster, and/or process more data without having to worry about architecture compatibility. Either the UNIX vendor delivered higher performance, expanded memory, and improved RAS, etc., or the customer would eventually migrate to a better platform. For decades, UNIX server vendors have provided the data center with continual, and timely, upgrade paths. This enables them to replace existing legacy environments with fully compatible servers that run the same operating environments on the same IT architectures with *higher performance*, *expanded memory capacity*, and *improved RAS* (reliability, availability, and serviceability) *characteristics*. In recent years, this cut-throat battle left three UNIX vendors focused on scale-up, UNIX applications dominating the landscape: IBM, with *AIX* running on *POWER* systems, HP with *HP-UX* on *Integrity*, and Sun, with *Solaris* running on *SPARC* systems. For some, expected upgrade plans have hit a major hurdle; however, with customers now assessing Oracle's acquisition of Sun, putting the future of SPARC under the microscope, and continued delays in the next generation of Intel's *Itanium* architecture (*Tukwila*), putting the future of Integrity in question, as well. Only IBM continues to provide the scale-up UNIX data center with a plan, and a predictable drumbeat, for the future. To learn more about IBM's forthcoming POWER7 and AIX7, please read on.

IN THIS ISSUE

- Enterprise Data Center Requirements .. 2
- A History of POWER 2
- POWER7 Takes Command 3
- Conclusion 4

Enterprise Data Center Requirements

In the past, the data center had simply needed to go faster and be less expensive. Now, however, the enterprise requires the data center to also *run smarter*. Medical facilities need to use IT to improve health care; the weather bureau needs to be able to make better predictions, especially as the next hurricane is headed toward the U.S. coastline. Furthermore, every data center needs to *improve its energy efficiency* with more intelligence being used to protect our natural resources.

In order to increase performance, lower cost, *and* run smarter, today's CIOs are looking to consolidate their mission-critical applications on fewer, faster servers and to virtualize those applications to make more efficient utilization of their data center resources, both human and natural. With the typical scale-out server operating at less than 20% utilization, **the data center has to change the operational paradigm to make better use of enterprise resources. To do this, the data center is looking for the server platform that has, among other requirements, the highest performance level in a virtualized environment.** Typically, that can be measured in terms of the number of *cores* available per processor, the number of *threads* available per core, the *bandwidth* on a CPU-to-CPU basis, the amount of *shared cache*, the amount of *memory* supported and the *memory bandwidth*, and the *I/O bandwidth* capability, not to mention the server *scalability* and *RAS functionality*.

In response to these requirements, Intel focused its efforts on improving its *Xeon* architecture to meet the data center needs for consolidation and virtualization. Specifically, it has announced the *Xeon 5500* (*Nehalem*) to achieve this functionality and reduce data center TCO. And, in fact, compared to the *Xeon 5400*, Nehalem can be considered an unqualified success, with significant improvements in database transaction count, integer and floating-point throughput, memory scalability, and memory bandwidth. Unfortunately for Xeon, the traditional enterprise data center has not been deploying an x86 environment for the past 40 years; it has been deploying applications based on mainframes, UNIX, or *AS/400* (*System i*). One of the reasons for that is reliability, with *Windows*-based servers exhibiting over 150 minutes of unscheduled outages per year, compared to, for example, only 15 minutes per year for *AIX*[1]-based *Power Systems* with less than one outage per year[2]. IBM's Power Systems are designed for RAS; designed, by the same team that defined what "mainframe-class" really means. Power Systems are deployed with RAS at every level: from the processor, to memory, to I/O, to the system itself.

Every UNIX-involved CIO needs to ask, "How does Xeon 5500 stack up against today's UNIX/RISC engines? How does it stack up against IBM's POWER architecture?" The reality is that IBM's *POWER6*-based platforms exceed the capability of servers built with the Xeon 5500, or for that matter Intel's *Itanium* microprocessor, in terms of database transactions, OLTP, integer throughput, and floating point capabilities. POWER6 also leads in terms of memory capacity per core, memory bandwidth per core, I/O bandwidth, and L2, L3 cache.

Over the past two decades, IBM has used its POWER architecture continually to improve the performance level of its Power Systems to improve the mission-critical application environment for the enterprise. Now, once again, IBM has announced further innovations in this architecture to enable more consolidation and better virtualization in order to encapsulate more workloads within each virtual machine, improving energy efficiency, and lowering TCO. As the leading world-class athlete in these arenas, when it arrives, ***POWER7* is expected to make its debut in the data center with performance and functionality to blow away the records set by POWER6.**

A History of POWER

IBM has always exhibited a regular development schedule for the POWER[3] architecture, from its inception in 1990, with improvements in performance and functionality at every step of the way. Over the past decade, *POWER4* added *LPAR*s (Logical Partitions) to the architecture in 2001; *POWER5* introduced *Advanced POWER Virtualization* in 2004; and *POWER6* provided *Live Partition Mobility* in 2007. Like a Swiss clock, every three years IBM has delivered innovations to the POWER architecture, including improved virtualization through *PowerVM*, along with corresponding improvements in the operating environment and systems management. **It must be noted here that IBM has a history of leadership in virtualization, dating back to the origins of the first hypervisor providing full virtualization – over 40 years ago – on the mainframe, which led to today's z/VM.**

[1] AIX is the IBM version of UNIX that runs on all POWER systems.

[2] Source: IBM.

[3] Performance Optimized With Enhanced RISC.

Every iteration of the POWER architecture has delivered improvements in the primary requirements for large scale computing for the data center. These include scalability, performance, RAS, and manageability. Equally important, however, are the improvements to the bottom line in *lowering the TCO* of the IT infrastructure. These improvements include:

- The **virtualization** of the AIX application suite – improving the utilization of Power servers up to 60 to 80% along with the capability for rapid deployment and dynamic change;
- **Energy Efficiency** to lower power consumption and reduce energy cost by up to 90%, delivering more work per watt with each new release of Power;
- Providing **Business Resiliency** for continuous availability and to dynamically respond to a changing business environment with no downtime or manual intervention, maintaining 99.997% uptime with only 15 minutes of unscheduled outage per year; and
- **Simplification of Management** through an ever-evolving IBM *Systems Director*, enabling control over deployments and energy usage.

Delivering a fully integrated solution gives IBM a tremendous advantage in consolidation and virtualization, enabling them to remove complexity from the IT environment and simplify data center operations before the platform is even deployed. Power Systems are designed, developed, and tested as an integrated server platform with a balanced systems design, enabling linear performance as core-count and utilization increase. Functionality is optimized across the layers of the system stack, including hardware, hypervisor, operating system (AIX, i, or Linux), middleware, and Systems Director management software. This enables Power Systems to deliver virtualization capabilities and attributes that are generally unmatched in an x86 server environment, including performance scalability, RAS, energy efficiency, and security. Simply put, the integrated combination of POWER architecture and PowerVM enables Power Systems to enjoy a higher consolidation ratio than any x86 server running VMware's *vSphere*, with up to 64 virtual CPUs and 4096GB of memory per VM under PowerVM compared with only 8 virtual CPUs with 255GB of memory with VMware *ESX 4.0*. For a more complete list of the emerging benefits of virtualization, see Exhibit 1, above.

Exhibit 1 – Emerging Benefits of Virtualization

- Reduced IT management costs
- Better Software investment protection
- Dynamic energy optimization
- Simplified High Availability and Disaster Recovery solutions
- Improved resource optimization
- Greater IT agility
- Turnkey packaged applications
- Improved security foundation

POWER7 Takes Command

With POWER6, IBM has the world's fastest processor at 5.0GHz, along with the only *tpmC*[4] result over six million transactions/minute. With POWER6, IBM introduced *Live Partition Mobility*, *Workload Partitions*, *Active Memory Sharing*, and multiple shared memory pools, along with *Active Energy Manager*: *EnergyScale*, and IBM *Systems Director VMControl* providing a consistent cross-platform management for physical and virtual environments. POWER7 will extend those capabilities to take a clear leadership role in the deployment of UNIX systems in the data center. POWER7 also will enable IBM to supercharge the System i platform, enabling them to deliver more performance and functionality for Linux on Power Systems[5].

With the looming probability of the sun setting on SPARC, and HP experiencing delays in the availability of the next version of Integrity due to delays in Tukwila, the quad-core Itanium processor[6], IBM, with POWER7 is in the right place at the right time. It has been reported as recently as this past June that Sun has officially killed off the 16-core *UltraSPARC-RK* processor, known internally as *Rock*, which has been in development for more than five years. In the meantime, Intel has indicated that RAS functionality from Itanium is being added to Nehalem-EX, perhaps signaling an end of life for Itanium.

With a similar portfolio of modular, scalable offerings as POWER6, POWER7 will be able to deliver a multi-core processor with a higher performance/watt, higher per core performance and more total systems performance. It is estimated that a Power7 System will have 2-to-5 times the total

[4] The Transaction Processing Council's performance benchmark.
[5] IBM has Power7 Systems running in the lab with AIX, IBM i, and Linux, all operational.
[6] It has been reported that the quad-core Itanium may not be available until 1Q10, two years later than originally planned.

system aggregate throughput plus a simplified, non-disruptive upgrade path through binary compatibility with POWER6 and Live Partition Mobility. Power7 Systems will possess greater levels of consolidation, up to 1000 virtual machines per system and better resource utilization than Power6 Systems due to dynamic, policy-based energy management and a single Systems Director console to manage the server and storage network. IBM has already announced the ability to upgrade a Power 595 system by simply replacing the processor books, including memory, and the system controllers. The 12X I/O drawers and GX adapters will migrate from Power6 to Power7 servers, preserving earlier investments.

Specifically, the POWER7 processor features eight processor cores to deliver four times the performance of POWER6 within the same power envelope. Each core will have 12 execution units, 4-way symmetrical multi-threading, and 256KB of L2 cache per core. Each processor will have 32MB of on chip eDRAM for L3 cache to be shared by the 32 threads. This eDRAM is nearly as fast as conventional SRAM, but requires far less space and standby power. There are dual DDR3 memory controllers, with four channels per controller, with a sustained memory bandwidth of 100GB/s and a scalability of up to 32 sockets, providing 360GB/s of SMP bandwidth. In addition, POWER7 includes a long list of RAS features to ensure application availability.

The major features of POWER7 Systems include linear scalability, PowerVM virtualization, both physical and virtual management, along with energy/thermal management. See Exhibit 2, above, for a list of virtualization features expected in POWER 7.

Conclusion

It is clear to even the most casual observer that IBM's Power Systems are in a league of their own today in providing the performance and advanced virtualization capabilities that are required for today's mission- and business-critical workloads. With over 40 years of virtualization experience and a commitment to integrate solutions across the systems stack, IBM continues to maintain a leadership position in the enterprise data center, even as others fade into IT history.

AIX, i, and Linux all run on the POWER architecture. These three operating environments represent almost half of all data center spending on servers[7]. In fact, IBM has reported that its Migration Factory is responsible for more than 1,750 migrations from Sun and HP platforms to IBM POWER systems over the past three years. Combined with the innovations in POWER7 and AIX, Systems Director and Tivoli, IBM cannot help but extend its lead in traditional IT environments and break new ground in Cloud Computing.

Exhibit 2 –
POWER Virtualization Today

- Fine-grained dynamic sharing of processors, memory, and I/O
- Optional dedicated resources
- Shared dedicated processors
- Extreme scalability and robustness
- Integrated firmware hypervisor
- Virtual I/O Servers layer
- Hardware enforced isolation
- LPARs and WPARs
- DLPAR and Processor Folding
- Capacity on Demand
- Partition Mobility

With better utilization of hardware, IBM is driving down the cost of software licensing, maintenance, floor space, and energy. It is reducing the TCO for Power Systems throughout the enterprise. Despite the turmoil in the rest of the UNIX market, IBM continues to outdo itself in delivering a high-performance, low-cost application solution to the data center. If your enterprise needs to improve the performance and lower the cost of your IT environment, IBM's Power Systems may be just the over-achieving world-record holder that you seek.

[7] Other spending is on Windows, Netware, z/OS, and others.

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.