THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2014012 May 13, 2014

IBM Opens Power Systems to a Much Broader Audience with More Performant and Cost-Effective POWER8 Servers

Analyst: David Reine

Management Summary

The more entrenched you are, the longer it takes you to change, or so it seems most of the time. You may have been with a bank or a stock broker or a mutual fund company for many years and you may see the world through that experience. You may know the ins and outs of doing business with them and probably there is a certain level of comfort that comes from that long-standing relationship. With time, you might have become a little more complacent than you would like to admit. It's not that the relationship hasn't been productive but now you want to expand into some new horizons. If you had little time invested, you'd probably be quick to change…*but now you see the light – change is necessary*.

That can be the lead in to a lot of storylines, but today I'm talking about IBM, in general, and IBM *Power Systems*, in particular. This isn't about right and wrong; it is about moving aggressively beyond one's traditional comfort zone. For decades, IBM's Power Systems (and its predecessor, the *RS6000*, before that) was a product line set in the *UNIX* world, where *AIX* was IBM's UNIX offering. Later, Power Systems ran both AIX and the *System i/i5* operating environments. More recently, *Linux* was added to the collection. For the longest time, the focus was on existing UNIX and System i customers, with UNIX not being limited to AIX environments, as IBM became very aggressive in the competitive UNIX space. Linux, which became a cross-platform initiative at IBM, was added originally to keep the UNIX and System i customers happy and to give them something else to run on these increasingly powerful servers with cycles to spare. More recently, say three or four years ago, IBM saw the light. The real threat was not other RISC/UNIX products from other vendors but the prevalent presumption that *x86* platforms were the right and proper platform for all things open. By this point in time, to many in the data center "open" meant "Linux". Clearly, IBM saw the light and decided to take its superior Power Systems platforms and compete directly for the hearts and minds of the open systems application developers. The fact that IBM is divesting its x86 server business, *System x*, to Lenovo, makes "Open on POWER" an even more important play for IBM.

Thus this is a two part story. The first is about IBM's OpenPOWER initiative and the OpenPOWER Foundation. The second is about a wide-ranging upgrade to POWER8 processors (and more) to the scale out line of Power Systems servers. Read on to learn more about both.

IN THIS ISSUE

The Changing Enterprise Data Center

In 2014, more than ever before, we see an even more rapid shift in the data center paradigm.

The need for people to "stay in touch" has led to a significant increase in the number of people using social media to keep track of their friends and family and to keep those same people appraised as to what they are doing. It has been estimated that 40% of the population socialize more on the Internet than they do face-to-face. The need to go to high school and college reunions has been replaced by posts to a Facebook wall. Many consumers have changed their shopping habits from "brick and mortar" to the Internet via their mobile devices, even when they are inside a physical store. Thus, brick and mortar merchants need to know that a prospective customer is in their store, what s/he is interested in, what s/he has bought previously, what other like-minded buyers might also have purchased, and what has been reported about the product and the merchant, etc. – all in real time. To do this in real time, the merchant and its data center must continue to innovate in terms of analytics, collected and available data, and innovative hardware and systems software capable of processing more data and more transactions more quickly, and also in an economical manner.

Currently, the innovative enterprise is amassing more and more unstructured data, often growing at a rate 15 times that of structured data. IBM has projected that there will be a 300 times growth of digital content between 2005 and 2020[1], and we are only halfway through that period now. The need for more storage and more powerful analytical solutions has never been greater. Procuring the right data center infrastructure certainly has become a critical business success factor.

Three major considerations tend to drive the infrastructure decisions of the enterprise data center.

- ***Is it open?*** – The infrastructure being deployed today must be on an open platform, it has to be supported by a large community, and it needs to be easily available to interested developers. Today's application developers prefer to work in open environments that are low-cost or no cost, readily available, and widely used.
- ***Is it optimized for Big Data?*** – Any server platform being deployed today for near real-time analytics must be designed and optimized for popular, widely-used Big Data tools and applications. Powerful, multi-core servers usually need more memory, more high-speed cache, and more I/O paths to accelerate the analytics processing required to make better and more informed decisions in short order.
- ***Does it provide superior value?*** – Value often is seen through the eyes of the beholder, but usually it boils down to cost per transaction or cost per analysis. Thus, this is a combination of hardware and software costs. The value of the server might be measured in terms of the number of cores needed to get the work done in a timely manner and the cost of the software needed to do the work. The usual presumption is that infrastructure will cost less if the server environment is open and seen as widely deployed and generally available.

The Holy Grail of Open Systems

The meaning of the phrase "Open Systems" always has been in the eye of the beholder. To some, if the underlying server platform isn't *x86*, it can't be open. To others, it is about availability and modifiability of source code. Some see universality (broad acceptance, applicability, and utility) and portability (without regard to specific server hardware) to be the defining criteria. Additionally, open tends to also mean "free of (or unencumbered by) costly operating systems and middleware". With that collected set of definitions, the path to open systems leads many to Linux. Those chartered with running enterprise applications want to run them in a widely available, cost-efficient yet robust virtualized environment, which usually leads to *KVM* or *OpenStack*. In the end, this tends to be more about open software than any specific hardware architecture.

To many, the very idea of an open server environment is synonymous with an *x86* architecture, whether from Intel or AMD or, possibly, some other widely-available architecture, such as *ARM*. Third-party software developers like writing code for these environments because they assume that these are the only acceptable environments. However, they do this without fully considering whether these are the best platforms for cost-effectively getting the work done in a short amount of time. In reality, if the requirement is to be open (think Linux, to start), to be Big Data focused, and to provide superior value per transaction or analysis, there is no actual requirement to be based on x86 or ARM or any other chip architecture.

[1] IBM presentation dated April 17, 2014.

Open Minds want OpenPOWER

The presumption that x86 architecture is required is being challenged by IBM's efforts to offer all of the above (without the x86 requirement, of course) on servers based on its *POWER8* processor (announced last month), especially on its *Power Systems* server platforms (also announced last month and which will be addressed later in this paper). Unfortunately, many in the IT community think of the *POWER*-processor-based environment as a proprietary architecture that runs *AIX* and *System i/i5* applications. IBM has been taking steps to correct that impression by also positioning POWER as an open system, built upon an OpenStack architecture, *Linux*, and, and other open interfaces, and by creating an open ecosystem[2]

All of these lead to *OpenPOWER* (a new focal point, i.e., the bright light) and the *OpenPOWER Foundation*, consisting of those vendors who together are making it happen. All of this will be discussed below. You need to keep an *open* mind, as that is a prerequisite for getting to your desired goals for open systems.

OpenPOWER

OpenPOWER is IBM's initiative to deliver an open systems operating environment[3] on the latest generation of servers based on the POWER architecture. To application developers, OpenPOWER largely is about a stack of operating software and middleware, because the hardware below the virtualization layer tends to be invisible to the developer – and that's the way they like it. However, OpenPOWER is more than that. It also is a consortium of interested parties that have banded together to accomplish two interrelated goals.

1. To encourage innovative uses of the POWER8 architecture – essentially to use the POWER8 processor and architecture as a superior building block (especially when compared to the x86); and
2. To assemble the components needed to complete and compliment OpenStack on that processor and architecture.

OpenPOWER Foundation

The Foundation is an open technical community based on the POWER architecture, enabling collaborative development and opportunity for member differentiation and industry growth. The goal of the OpenPOWER Foundation is to create an open ecosystem centered on the POWER architecture and to share expertise, investment, and server-class intellectual property to serve the evolving needs of customers and industry.[4]

The Foundation is positioning and advocating IBM POWER-based open servers and solutions as a better alternative to x86-based and ARM-based servers and open systems stacks. It is doing this by simplifying system design across the hardware/software stack and implementing a Little Endian[5] Linux, originally designed for x86 architectures, to ease the migration of applications to IBM Power Systems and other servers based on the POWER Architecture, and to build a complete, flexible, and enterprise-class ecosystem around that architecture.

IBM has led the effort in the development of an open technology platform with an industry consortium focused on innovation. This platform has been built with an OpenStack architecture across both hardware and software, including *KVM* for cross-platform universality, *CAPI*[6] for improved integration between POWER8 and the open development community and, of course, the portability of Linux. Its purpose is to deploy a high-performance, open systems, viable alternative to the x86 architecture, that will leverage open collaboration across the stack and revolutionize the way that IT is deployed and delivered, thus providing more value and new, superior choices to the enterprise data center.

Several industry leaders, including Google, Mellanox, Micron, Nvidia, Samsung, PowerCore, and Tyan have joined with IBM in this effort to propagate OpenPOWER throughout the enterprise data center. This includes members driving everything from technology and microprocessor fabrication to system design, to firmware and software development, including I/O, memory, networking, and storage.

---

[2] In support of this initiative, IBM last year committed $1 billion (USD) in new Linux and other open source technologies for IBM's Power Systems servers. Major investments include new products, a growing network of five Power Systems Linux Centers around the world and the Power Development Platform, a no-charge development cloud for developers to test and port x86-based applications to the Power platform.

[3] *AIX* is not involved.

[4] For more information, see http://openpowerfoundation.org/.

[5] Big Endian and Little Endian are different approaches to representing bytes in words, more specifically, how the bytes are ordered within the word. Little Endian systems are those in which the least significant byte is stored in the smallest address. See http://en.wikipedia.org/wiki/Little_Endian.

[6] Coherent Accelerator Processor Interface.

Exhibit 1 — The OpenPOWER Community

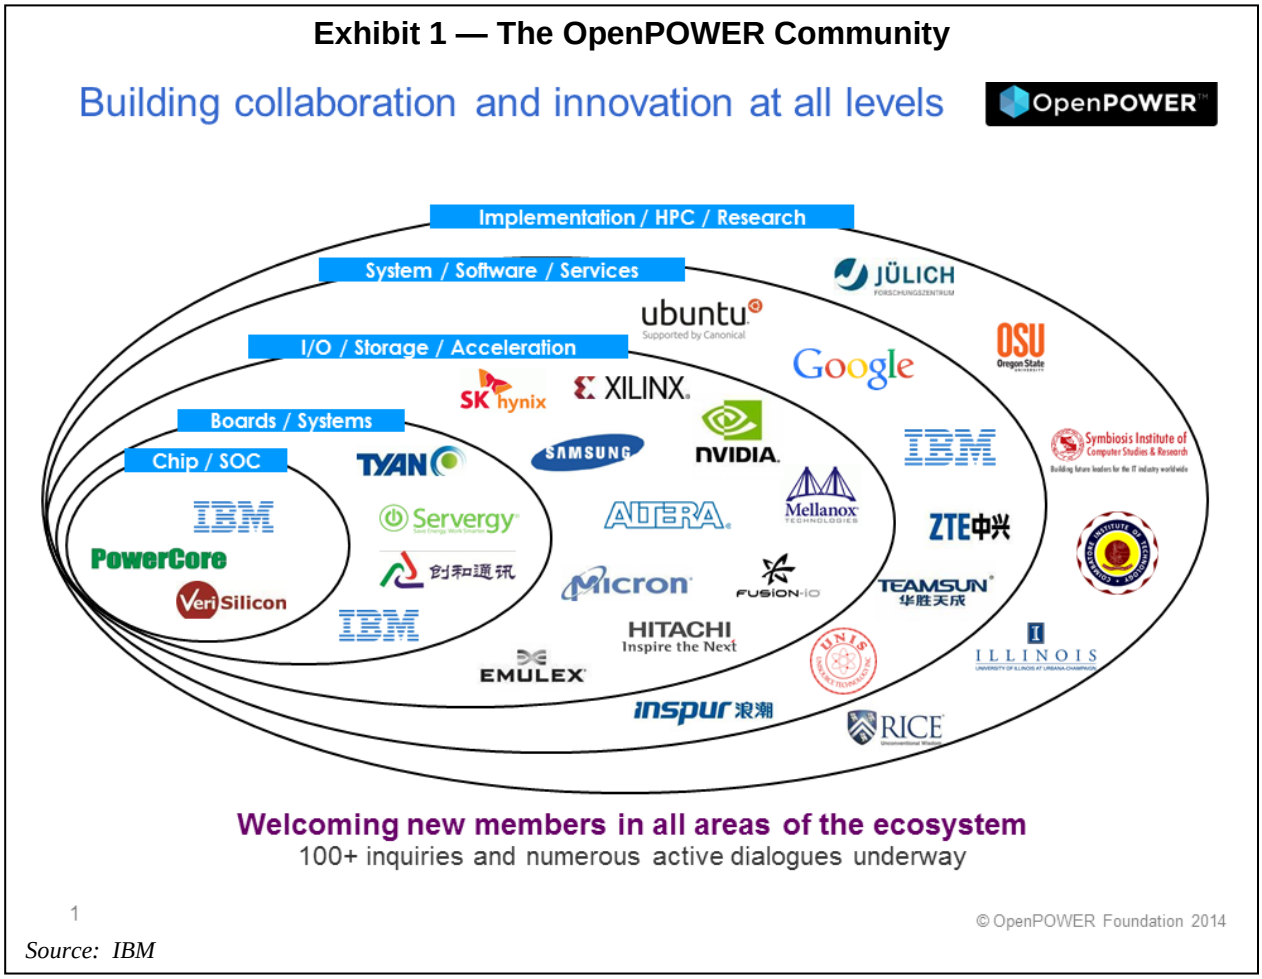

Source: IBM

As shown in Exhibit 1 above, there's a lot more to this story than just IBM. Many companies are involved, across a broad spectrum of technologies. Of course, each has its own reasons, but they all share one goal – to make the whole greater than the sum of the parts. The Foundation plans to do this by opening the architecture to give the industry the ability to innovate across the full hardware and software stack. It wants to achieve this by:

- Simplifying system design with a superior alternative to x86 alternatives.
- Including SOC design, bus specifications, reference designs, and firmware, operating system and hypervisor as open source
- Incorporating Little Endian Linux to ease the migration of software to POWER (primarily from x86, which also is Little Endian). By sharing this approach with the physical addressing scheme used by x86 processors, life is made simpler for programmers.
- Driving an expansion of the enterprise-class hardware and software stack for the data center.
- Building a complete ecosystem to provide customers with the flexibility to build servers best suited to the POWER architecture.

Open Application Development

Of greatest interest to application developers is the openness of the operating environment. This might be better seen in terms of cloud infrastructure. Usually, applications developed to run in a cloud have little visibility of the hardware that is underneath. This appeals to application developers, because it allows them to write code that is independent of the infrastructure underneath. This transparency is enabled via OpenPOWER in several ways.

- ***Linux*** – IBM has enhanced the POWER8 environment with a broad Linux ecosystem to provide the simple migration of <u>all</u> Linux workloads to the POWER architecture. No changes are required for most well-written Linux applications, i.e., those written in scripting or interpretive languages. A simple recompile is required for an estimated 95%

of Linux applications written in C or C++. OpenPOWER is qualified with all of the most common Linux implementations from *Red Hat*, *SUSE*, and *Ubuntu*.

- ***OpenStack*** – The OpenStack mission is to develop an Open Source cloud computing platform that will meet the needs of both public and private clouds regardless of size and to provide ease of implementation and massive scalability.[7] OpenStack enables a new private cloud offering that significantly speeds and simplifies managing an enterprise-grade cloud. Enterprises now have a core set of open source-based technologies to build enterprise-class cloud services that can be ported across hybrid cloud environments. This will ensure that innovation in cloud computing is not hampered by locking businesses into proprietary islands of insecure and difficult-to-manage offerings.
- ***KVM*** –The *Kernel-based Virtual Machine* (*KVM*) is an open source hypervisor that provides enterprise-class performance, scalability and security to run *Windows* and Linux workloads. KVM provides organizations a cost-effective alternative to other x86 hypervisors, and enables a lower-cost, more scalable, and a smarter choice for open access to the cloud.

The IBM POWER8 Architecture

As was indicated at the outset, this report has two storylines: OpenPOWER (just covered) and IBM Power Systems based on the latest POWER8 processors. Now, let's look at the POWER8 processor and architecture.

In the forefront of the move to OpenPOWER is POWER8. It is exactly what it sounds like, the eighth generation of IBM's POWER architecture, continuing the performance and reliability of IBM's successful *RS/6000* systems. Dating back to 1990, POWER delivers the high performance and reliability that IBM's mainframe customers had come to expect. The POWER8 processor has been designed to address all of the Big Data needs of enterprise data centers across the globe, including compute, memory, I/O, and storage, and using a 22nm process rather than the 32nm process used for *POWER7+*. An innovative core architecture enables POWER8 to double most resources from the previous generation, while enabling the data center to deploy from eight to twelve cores per socket, with twice as many threads as POWER7 cores. POWER8 processors have 512KB L2 SRAM per core with 96MB of eDRAM in a shared L3 cache and up to 128MB of L4 eDRAM, all significantly greater than in POWER7.

Compute capability

POWER8 is designed as a massively multi-threaded processor, capable of handling 96 hardware threads simultaneously. POWER8 has up to 12 flexible, fast cores, 50% more than POWER7, that have been designed to run analytics algorithms better than an x86 platform[8]. Each processor has four times the number of threads per core (in comparison to an Intel *Xeon Ivy Bridge* x86 architecture), with a 30% higher clock (operating) frequency.[9] Using GPUs from OpenPOWER partner Nvidia, IBM POWER8-based Power Systems can deliver three times the query acceleration on *DB2* running on the POWER architecture and Linux[10]. IBM has reported that for most workloads, POWER8 will perform two to three times faster than POWER7.

Memory

POWER8 architecture also has been designed for more memory to support a larger working data set to maximize business insight. POWER8 delivers four times the memory bandwidth of an Ivy Bridge x86 platform with 30% higher memory capacity. .

I/O Network

POWER8 has a larger I/O bandwidth in order to support more data in motion. In fact, POWER8 has five times the I/O bandwidth of IBM's POWER7 Systems. This far outpaces the x86 capability. According to IBM, a POWER8-based distributed store accelerated by RDMA can deliver ten times the throughput and seven times the latency for in-memory stores.

CAPI

CAPI is an innovative gateway for external communication allowing the industry as a whole to directly integrate partner value into POWER8 chip protocols for Big Data applications, with up to a 20 times improvement in access efficiency. This provided cache-coherent access to DRAM, allowing many enterprise applications

[7] See http://www.openstack.org/.

[8] Please note that all comparisons with x86-based systems reference various versions of Intel's *Sandy Bridge* architecture.

[9] Higher is better.

[10] In addition to Linux, POWER8 supports IBM *AIX* and *IBM i5OS*.

Exhibit 2 —
A Broad Range of Power Systems Servers based on POWER8 Technology

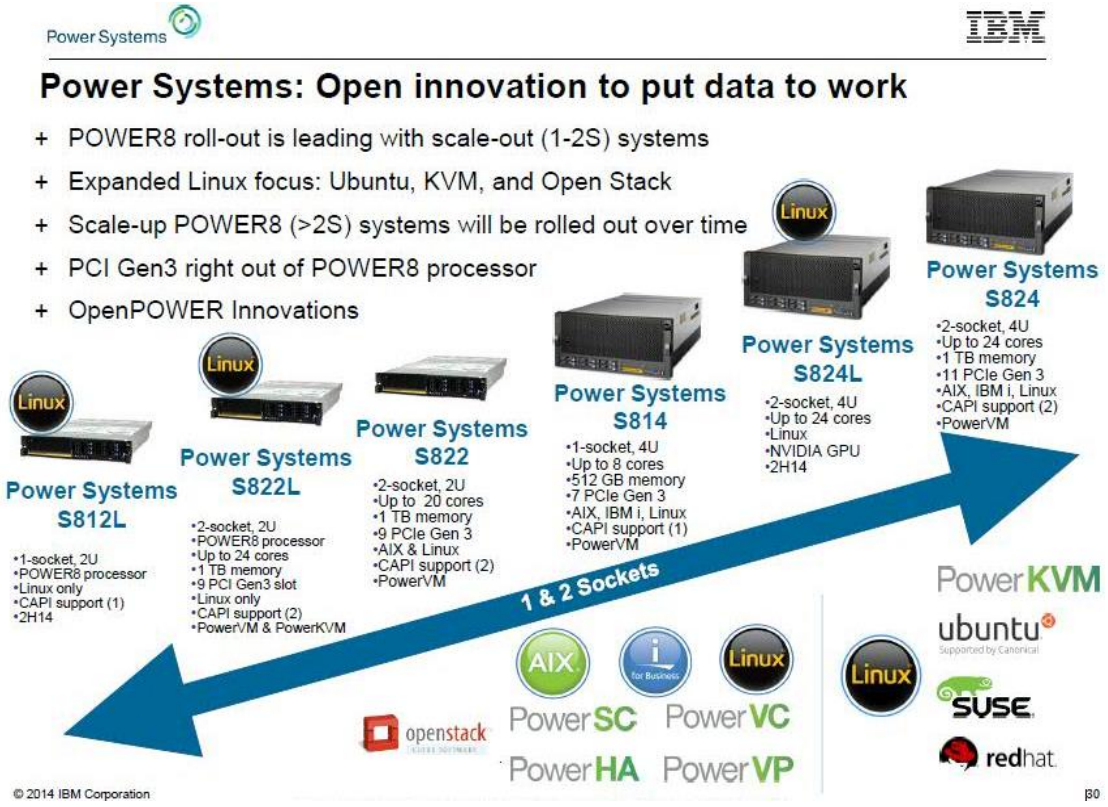

Source: IBM

to bypass some thick protocol layers by using *POWER8* chip protocols, thus significantly accelerating in-memory access. Along with a *Dynamic Balanced Memory Architecture (DBMA)* and an accelerator to support virtual addressing that can work with the same memory addressing scheme that the POWER processors use, POWER8 can eliminate up to 97% of the instruction data path between flash and the application, thus delivering a 2.6-times improvement over x86 in internal data flow, a 4-times improvement in memory bandwidth, and a 3.2-times gain in cache capacity. This enables the lowest possible overhead and latency for highly referenced data from a shared data space, enabling workload optimized systems for analytics, transaction processing, and Big Data in the cloud solutions.

IBM Power8 Systems

With POWER8 technology, IBM has created the next generation of Power Systems. These POWER8-based servers can deliver compelling application performance with extreme flash optimization, with an easy upgrade path from *POWER6* and *POWER7* servers. IBM has not only configured a complete family of Power Systems (see Exhibit 2 above) designed with the POWER8 processor; it has, in fact, made available two families of one- and two-socket scale-out servers, one designed specifically for Linux and the other designed for all of their operating systems, Linux, AIX, and IBM i.

What does all of this mean for the enterprise data center and the applications that drive it? For one thing, according to IBM, POWER8 Systems with in-memory data and with *DB2* and *BLU* acceleration can deliver business insights up to 50 times faster than with servers based on Intel Ivy Bridge and a traditional database. This is a result of POWER8 being the first processor designed specifically for Big Data with massive parallelism and bandwidth to accelerate delivery of real-time results. This means that enterprise executives often can receive most complex reports in less than 1/100th of the time.[11]

For another, POWER8 Systems are ready for deployment in the cloud, whether that is a

[11] According to IBM.

private, public, or hybrid cloud. Designed with a single, secure user interface built on OpenStack APIs, these systems can manage whatever cloud deployment makes sense for the enterprise. This is what application developers want and data centers need.

Additionally, POWER8 Systems provide up to 26TB of internal flash with automated tiering of locally attached HDDs and SSDs. Using flash memory, they provide the fastest possible access to volumes of high priority stored data. With the activation of compression acceleration, POWER8 Systems enable a four times storage capacity improvement with a two times increase in effective bandwidth.

IBM Power Systems for POWER8 with Linux

IBM offers three models in the POWER8 family that have been specifically configured only for Linux.

- ***The Power Systems S812L***, a one-socket, 2U server that will be available in the second half of 2014;
- ***The Power Systems S822L*** a two-socket, 2U server with up to 24 cores and 1TB of memory and nine PCI Gen3 slots; and
- ***The Power Systems S824L***, a two-socket, 4U server with up to 24 cores.

All of these models have been priced to compete competitively with Linux on the x86, delivering 20-30% better price/performance when compared to the same stack on Intel Ivy Bridge-based servers, according to IBM. All of the models will run the Red Hat, SUSE, and Ubuntu Linux implementations in an OpenStack environment. See Exhibit 3, at the top of the next column, for a description of additional new Power8 Systems models capable of running AIX, IBM i, and Linux.

Conclusion

IBM clearly is driving innovation across its entire scale-out Power Systems portfolio with its POWER8 processors and architecture and open collaboration with many OpenPOWER partners. Together, they have strategically endowed POWER8 with all of the forces driving server acquisitions today – Big Data, open system software with a focus on Linux and OpenStack, and cloud computing.

IBM has increased core throughput for POWER8 thus taking a significant step forward in per socket performance and enabling a more robust multi-socket scaling. Along with the integrated access to memory from CAPI, IBM has accelerated business processing with faster time-to-insight for critical business decisions and up to 50 times faster in a Big Data environment. With the POWER8 changes to the architecture and next generation advances to Power Systems, the enterprise data center can achieve a 24:1 server footprint reduction in a *noSQL* environment.[12] In addition, with a non-disruptive application mobility from POWER6 and POWER7 Power Systems, the data center can protect prior enterprise investments in IBM Power Systems solutions.

Along with the continued support from its partners in the OpenPOWER Foundation, it certainly appears that IBM is committed to delivering more value to the enterprise data center for not only the existing application set, but for new workloads, as well, all while creating faster time- to-business insight. If your executives are looking for everything they need to make their business decisions faster, look into IBM's new Power Systems based on POWER 8 technology. They may have the answers that your executives are seeking.

Exhibit 3 — Additional Power8 Systems for AIX, IBM i and Linux

IBM now offers three models of the POWER8 family to provide the performance and reliability expected from a system capable of running AIX, IBM i, and/or Linux at or below the price/performance of an x86 system.

- ***The Power Systems S822*** is a two-socket, 2U server, with up to 10 cores per socket. It can support up to 1TB of memory.
- ***The Power Systems S814*** is a one-socket, 4U server with up to eight cores and 512GB of memory and seven PCIe Gen3 slots.
- ***The Power Systems S824*** is a two-socket, 4U server with up to 1TB of memory and 11 PCIe slots.

Source: IBM

[12] According to IBM.

About The Clipper Group, Inc.

The Clipper Group, Inc., now in its twenty-second year, is an independent publishing and consulting firm specializing in acquisition decisions and strategic advice regarding complex, enterprise-class information technologies. Our team of industry professionals averages more than 40 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

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.