THE CLIPPER GROUP
Navigator™

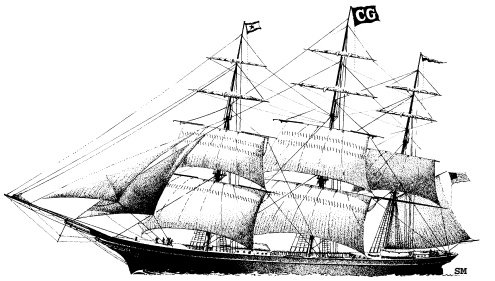

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2008059 October 29, 2008

IBM i – Using a Common Platform to Reduce TCO with Higher Performance and Lower Energy Consumption

Analyst: David Reine

Management Summary

If you are looking for the ideal model for efficiency in performance and energy consumption, one need look no further than the transportation industry for a lesson in lowering costs and improving profit. The flexibility of railroads around the world provides a lesson in how to adapt to the environment with a common platform. The railroad engine, or locomotive, is designed to pull a tremendous amount of weight, regardless of content, on tracks that are uniform, at least in most geographies. It does not matter if the train consists of freight cars or passenger cars; the locomotive applies the appropriate amount of power to move the contents from *Point A* to *Point B*. What's more, a train can consist of a heterogeneous mix of cars: freight, passengers, baggage, even cars the automobiles of the passengers. If you want to transport more than a single locomotive can handle, add a second engine, and double the workload. The locomotive does not ask what it is pulling; it just pulls or pushes.

Similarly, the enterprise data center has a variable workload of legacy applications that need to run on a daily basis. The CIO has deployed these mission- and business-critical applications with a variety of operating environments, always matching the application to the environment best suited for it. In most cases, the enterprise selects the *application* that meets its needs first, and then implements the operating environment required. Unfortunately, in the past, the IT staff has had to deploy a rather wide, heterogeneous mix of platforms in order to provide the operating environment with the proper support system to run the data center efficiently. From an historical perspective, the enterprise has deployed *UNIX* systems for some applications, and applications written for the *AS/400* environment for others. Over the years, the IT staff implemented UNIX applications more often than not on AIX servers from IBM, based on what was IBM's *System p*. The AS/400 applications have a long and storied past, too long to review here. Suffice it to say, yesterday's AS/400 applications are still running today on IBM's *i5/OS* which now finds its home on IBM *i*. Today, many data centers have added a third component to this mix, applications based upon an open systems version of UNIX, *Linux*. In fact, many data centers have deployed Linux applications on a heterogeneous mix of x86 platforms using either the Intel *Xeon* microprocessor or AMD's *Opteron*. All of these different servers require the enterprise to support a complex infrastructure of diverse platforms, increasing the management requirements and, therefore, the total cost of ownership, or TCO. What can the data center do?

In an effort to simplify the infrastructure, and the life of the CIO, IBM has consolidated all three of these environments under a common *POWER6* architecture. Today, an IBM *POWER System* enables the data center to drive all three environments, much like our railroad system, on a common platform, with a single set of tracks, to simplify the IT environment. To learn more about POWER systems, please read on.

IN THIS ISSUE

A Complex Heterogeneous Data Center

Today's enterprise data center is a heterogeneous mix of mission- and business-critical applications running on a diverse set of operating platforms. As a result, the IT staff must maintain and support not only multiple operating environments, deployed over a variety of server platforms, spanning many years and processor generations, but multiple teams of support personnel, as well. This server sprawl has created a great deal of complexity in the data center, and complexity always adds cost to the enterprise operational budget. Simplifying the infrastructure, and reducing the TCO of the data center, has thus become a major focus of the enterprise executive team.

Some enterprises have looked at changing the IT environment, and migrating to a scale-out architecture using a network or grid of x86 servers to simplify the infrastructure. Unfortunately, when deployed with a single application resident, these servers only utilize a fraction of the CPU resources available, in many cases less than 20% of server compute cycles. This wastes not only valuable IT resources, but natural energy resources as well. More importantly, these enterprise data centers are *application-centric*, not platform-centric. The CIO is most concerned with retaining the investment that the enterprise has made in these mission- and business-critical applications, along with the investment made in support personnel and management software. Now, can the data center reduce the TCO without affecting the performance and reliability of the IT infrastructure?

In order to reverse this trend, simplify the IT architecture and lower data center TCO, the IT staff must change the paradigm of the enterprise data center, focusing on:

- The consolidation of mission-critical server and storage platforms;
- The virtualization of those platforms to improve server and storage utilization;
- The deployment of a storage area network to facilitate that consolidation; and
- The implementation of programs in support of green IT, to lower energy consumption and help to reduce the TCO of the IT environment.

Legacy data centers with a preponderance of mission-critical applications based on IBM's AS/400 and i5/OS operating systems, along with UNIX and Linux applications, have had the opportunity this year to review, and consolidate on, the latest platforms from IBM Power Systems, designed to host all legacy applications written to either the AS/400 or AIX operating environment.

POWER Systems Integration

Earlier this year, in April, IBM introduced a new line of Power Systems to provide small and medium business (SMB) data centers using IBM *System i* and IBM *System p* platforms with a unified line of servers to increase application choice, reduce energy consumption, and reduce administration and maintenance costs. Furthermore, it also provided these customers with a simplified and reduced pricing structure. Configured with the latest dual-core *POWER6* processors, these Power Systems provided the data center with a powerful tool for the consolidation of *IBM i*[1], AIX[2], and Linux applications on a single, energy-efficient Power platform. With POWER's *Capacity on Demand*, the data center can increase processor and memory resources as needed, without disruption, and then turn it off when the increased workload subsides. Power Systems support AIX V5.3 or later, IBM i 6.1 or later, *SUSE Linux Enterprise Server 10 for POWER* or later, and *Red Hat Enterprise Linux 4.5 for POWER* or later. This enables the data center to choose the operating system that best suits its business needs.

In fact, with these new high-performance POWER6 microprocessors enabling them, data centers could migrate their existing mission-critical System i applications to the new servers, providing simultaneous support for all three environments, increasing performance and energy efficiency, using the same IBM i applications and operating system that they have relied upon for the past two decades. At the same time, they gain access to the complete library of AIX and Linux applications that have been running on System p servers. In addition, the System i data center also gained access to IBM's *PowerVM* virtualization technology, enabling the IT staff to create up to 80 virtual partitions on a single

[1] AS/400 or i5/OS.
[2] UNIX.

platform, deskside, rack, or blade. With PowerVM, the data center can create a *scale-in* architecture, enabling a higher utilization of server resources. With *POWER6 EnergyScale* technology, the data center can provide the System i user with the same, advanced energy control features that the System p user has had to help manage energy costs. This announcement simplified data center management for both the enterprise and SMB.

POWER6 Enhancements

With the initial announcement of Power Systems, IBM enabled SMBs with *Express* offerings for *BladeCenter* and the low end of the *Power 520* Express and *Power 550 Express*. Now, IBM has extended that offering, making it more attractive for the enterprise, or departments within the enterprise, to access the consolidation and virtualization advantages of the converged Power System. IBM has announced new POWER6 based models to support up to 16 POWER6 cores in a single scale-up[3], scale-in system, along with updated virtualization and systems management software to enable the data center to improve performance while maximizing return on their IT investment. This also provides additional head room for SMB customers, protecting their investment in Power Systems.

In addition to the *JS12 Blade*, *Power520 Express* and the *Power 550* Express announced in April, IBM has now introduced the *Power 560 Express*, along with upgrades for the Power 520 and Power 550.

Power 560 Express

IBM has introduced the *Power 560 Express* server for consolidation of AIX, IBM i, Linux for Power, and x86 Linux workloads for mid-size-to-large database and application platforms, providing more headroom than the Power 550. The Power 560, with up to 16 64-bit POWER6 cores running at 3.6GHz, provides the performance needed to be an ideal virtualization engine. Combined with IBM's PowerVM technology, the Power 560 can aggregate and manage all application resources while helping to simplify and optimize the IT infrastructure and reduce server sprawl. The Power 560 also carries an extensive portfolio of qualified, proven solutions from the AIX, IBM i, and Linux catalogs.

The Power 560 is configurable as a single building block in a scalable 4U rack-mounted chassis, with the first node supporting either four or eight POWER6 cores. With a second chassis, the 560 can scale the system up to 16 cores and up to 384GB of memory, while leveraging the Power System's symmetrical multiprocessing (SMP) architecture for application consolidation, reducing overall energy consumption, server footprint in the data center, and software licensing costs. Each processor core has 4MB of L2 cache, with 32MB of L3 cache shared between each pair of cores. Each building block supports up to six 3.5", 450GB SAS drives, for 2.7TB of internal storage per chassis.

At the same time, the IT staff can utilize IBM's PowerVM technology to achieve the scale-in quality necessary to improve overall system utilization while protecting the enterprise investment in IBM i and AIX resources. This technology includes the standard *POWER Hypervisor* with *LPAR*, *Dynamic LPAR*, and *Virtual LAN* technology. The optional *PowerVM Standard Edition* provides the Power System with micropartitioning for up to 10 micropartitions per processor, along with multiple shared processor pools and virtual I/O server capability.

As with all Power Systems, the 560 has the reliability, availability, and serviceability features required of a mission-critical server to deliver near-continuous application availability. These features include the capability to recover from intermittent errors or failover to redundant components. The Power 560 can also detect and report failures, or pending failures, and includes the ability to initiate action automatically to effect error correction, repair, or component replacement. Additionally, the Power 560 can redirect workloads to alternative processors, without disruption to application execution.

Power 560 Performance

The Power 560 has outstanding performance for a System i environment, with a *CPW*[4] rating of 14,100 for four cores, 27,600 for eight cores, and 48,500 for 16 cores, showing near-linear scalability. This represents, about, a 70% improvement over System i POWER5 models,

[3] See the issue of Clipper Notes dated September 23, 2008, entitled *Perceiving the Dark Side of the Moon – Knowing When Scale-up Computing Makes Sense*, and available at http://www.clipper.com/research/TCG2008048.pdf.

[4] Commercial Processing Workload

and a 30% improvement over System i POWER5+ servers. In compute intensive environments, it is better to use a more commodity measurement, such as the integer and floating point benchmarks from *SPEC*[5]. With the SPEC ratings, we can get a better view as to how the Power 560 compares to other scale-up platforms, such as the HP *rx7640*, based upon the Intel *Itanium* microprocessor, and the Sun *M5000*, based upon *SPARC*.

Preliminary results for the Power 560 with 8 processors and 16 cores indicate a *SPECint_rate* value of 363 and a *SPECfp_rate* value of 263. These figures compare quite favorably against both HP and Sun. The *HP rx7640* with 8 *Itanium* processors and 16 cores has a SPECint_rate value of 201 and a SPECfp_rate value of 174. This represents almost an 80% edge in integer math and a 50% advantage in floating point calculations. The Sun SPARC M5000, *with 32 cores*, has a SPECint_rate value of 264 and a SPECfp_rate value of 223. This represents a 38% edge in .integer math and an 18% edge in floating point arithmetic. In either case, it is no contest in pure performance with POWER6 having a distinct advantage over both Itanium and SPARC. As would be expected, the delta between either Xeon or Opteron and POWER6 is even greater.

As a result of the performance levels of the POWER6 microprocessor, the data center can do more work with fewer processors, thus reducing the amount of infrastructure required, reducing floor space and energy consumption, possibly reducing software licensing costs, and lowering the TCO of the data center.

Power System Enhancements

In addition to the introduction of the Power 560 Express, IBM has also improved the performance and capacity of several previously-announced platforms in order to provide the System i data center room for business growth or to add new Linux or UNIX applications.. The Power 520 Express now has an option for a 4-core deployment, while the Power 550 has doubled its capacity to an 8-core server. IBM has also extended the capability for full integration support to *BladeCenter*, with the preinstallation of both IBM i and AIX on the *JS12* blade server. IBM has also enhanced the Power Systems management software suite with new versions of IBM *System Director*, *Active Energy Manager*, PowerVM *Active Memory Sharing*, and *PowerHA* for AIX. (See Exhibit 1, above.)

Exhibit 1 –
Enhanced Power Systems Software

- **IBM Systems Director** – Helps the data center deploy, monitor, analyze, optimize, and update Power Systems running any combination of IBM i, AIX, and Linux;
- **PowerVM Active Memory Sharing** – Helps the IT staff improve memory utilization by pooling resources between partitions;
- **Active Energy Manager** – Utilizes advanced energy control options to boost performance per watt by slowing processor clock speed or putting processors in "nap" mode when not in use and enables the IT staff to set an energy cap for a single POWER6 processor or a pool of POWER6 processors; and
- **PowerHA for AIX** has been enhanced with asynchronous GLVM support to enable geographic distribution of systems to improve business continuance and disaster recovery.

Source: IBM

Conclusion

With server sprawl rampant throughput the data center, causing an increase in the TCO of the data center because of under-utilized servers and wasted energy, **consolidation is not an option, it is a requirement**. For any enterprise with legacy applications under AS/400 or i5/OS, along with UNIX applications, IBM's Power Systems become a no-brainer. Deployment of a Power System enables the data center to reduce the complexity of the mission-critical operating environment, improve resource utilization, free up floor space, reduce energy consumption, and by the way, gain access to an impressive catalog of Linux applications, as well. IBM's Power System may just be the way to improve your bottom line, too.

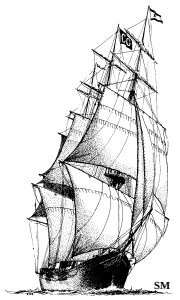

[5] Standard Performance Evaluation Corporation.

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.