THE CLIPPER GROUP
Navigator™

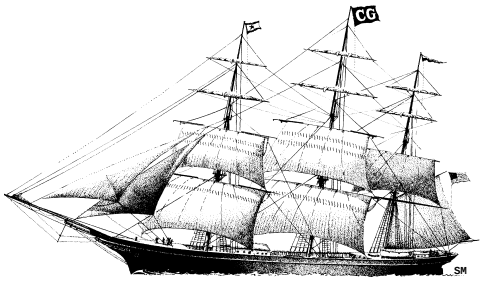

Navigating Information Technology Horizons℠

Published Since 1993 Report #TCG2006076 August 24, 2006

IBM's System p — Pathway to Consolidation

Analyst: David Reine

Management Summary

As mid-summer arrives and the dog days of August are upon us, we can do little about the weather except watch the temperature climb into triple digits and try to find relief from the heat. We can sit in our homes and crank up the air conditioning, while the electric meter spins so fast that you can develop vertigo from watching it, *and go broke from paying the utility bill as the cost of energy continues to climb.* We can do the laundry at night when electric energy demands are less or turn off some lights, which are growing dim anyway from the demands that we put upon our energy providers, but measures that are more significant either are or will be called for. We need to replace older appliances, the refrigerator, dryer, TV, etc., that waste watts, with more energy efficient models, improve the insulation in the exterior walls in our homes, perhaps even replace those windows that allow heat and cooling to pass right through with more efficient triple-pane units. It would be nice if we could increase the amount of electricity that arrives through the fuse box or circuit breaker panel, but we cannot. For most of us, we have to learn to reduce the amount of energy that our homes consume. Not only will this conserve our natural resources, but it will also reduce the total cost of ownership (TCO) of our homes.

The very same problems exist in every office building, in every city, except magnified because the demands occur in the middle of the day, when the demand for energy is at its peak. Enterprise Information Technology (IT) infrastructure uses, and wastes, electricity, not by the *watt*, but by hundreds or thousands of k*ilowatts per hour*. When the lights dim in the office, more than just a TV is affected. Any one of a vast array of under-utilized servers (like home appliances) that drive the data center may crash, leaving the enterprise isolated from both employees and customers, costing the enterprise huge amounts of money. The CIO must improve IT efficiency to lessen the total cost of ownership (TCO) of technology and to decrease enterprise dependence on utilities.

One course of action available is to *consolidate enterprise resources* on more powerful and price/performant platforms based on today's energy efficient architectures. One such solution is the *System p* family based upon the *POWER* architecture from IBM. Already the highest-rated transaction engine, IBM has improved not only the high-end of the *p5* series, but the entire family, with the latest processor innovations in the p5+. However, physics alone will not provide a business solution for your enterprise. That requires a support team of experts who know how to frame, plan, and design a solution, as well as a staff ready to assist in the implementation, in order to mitigate the risk to your data center. IBM has put such a team together. To learn how IBM and System p can lower the TCO of your data center, please read on.

IN THIS ISSUE

Data Center Server Utilization

The key issues facing the data center in every enterprise focus around *server utilization, growth, physical and logical consolidation of enterprise resources, energy cost, and the continued use of legacy applications*. The CIO needs to reduce the TCO of IT resources, from acquisition costs, to operating expenses such as data center staffing, and the energy required to run and cool the data center. Existing scale-out infrastructures consist of a plethora of aging, underutilized mono- and dual-processor nodes clustered together. Some larger enterprises have implemented a scale-up environment, using a large number of processors, usually 16 to 32, or more, networked together in a symmetrical multiprocessing (SMP) architecture. Clearly, these SMP systems utilize the available CPU cycles better through a virtualized operating environment, however, their operating costs, from maintenance to energy, continue to rise due to their age and the nature of their single-core, high-clock-speed architecture.

Many enterprises are now considering a move to scalable, general-purpose, commodity servers in order to consolidate and improve the utilization of server infrastructure. They are *perceived* to have a lower purchase cost and a higher price/performance ratio. Unfortunately, these CISC[1] systems lack the on-demand scalability available in some legacy RISC[2] systems, such as IBM's *System p*. CIOs can take advantage of a new warranty and lower maintenance costs to offset acquisition costs and lower data center TCO. In fact, many data centers are still dependent on mission-critical applications that require the higher reliability, availability, and serviceability (RAS) traits available in long-standing operating environments, such as *AIX* and *HP-UX*, and need to continue using them to facilitate growth. The availability of *Linux*, however, is essential in order to access a wealth of commodity applications appearing from the open systems community. Virtualization techniques enable improved utilization by allowing the enterprise to run multiple mission-critical *and* commodity environments at the same time. Enterprises with existing AIX applications can virtualize their newer POWER platforms and share processing cycles between their AIX and Linux applications, improving server utilization and reducing operating costs.

Many installed open x86 servers are running at less than 20% efficiency, wasting critical natural resources and leaving significant unused compute time. Many remote departments and smaller enterprises are now restricted from additional growth, while enterprise data centers are reaching their limit for floor space and available power to the facilities. The CIO needs to consolidate those servers to improve CPU utilization and relocate the server resources in order to reduce the administrative overhead and staffing. Increasing the density and utilization of server and storage resources will reduce the power requirements, and simplify the application environment to facilitate growth.

Reducing energy requirements for running and cooling the data center efficiently is of paramount importance. With the cost of energy rising to previously unheard of limits, it becomes necessary for every data center to increase the performance per watt in order to reduce the recurring costs of data center operation. With one eye on the bottom line and the other on the horizon, the CIO must protect the enterprise's investment with a scalable platform, capable of adding, and upgrading, processors and memory without the aid of a forklift.

In order to avoid the addition of staff and minimize the costs of training, maintenance, and integration, many enterprises are looking to standardize on a single platform architecture to achieve the economies of scale desired, much the same way that a company such as Southwest Airlines has standardized on a single airplane model, the *Boeing 737*. A standard operating environment enables the enterprise to implement a uniform, scalable processing model, along with scalable I/O, to take advantage of future architectural advances. Any new platform must be able to protect the investment in storage solutions already installed, be it DAS, SAN, or NAS, as well as have access to the latest in low-cost, high capacity disk architecture. This will enable the data center to maximize the utilization of business-critical data for applications that take advantage of D2D storage techniques, such as backup/ recovery and archiving. Standardizing the consolidation on one platform architecture also enables the enterprise to simplify network manageability, reducing administrative overhead. The question is: *Which platform?*

IBM and System p

IBM has a long history with RISC architectures and *POWER* CPUs, dating back to the first RS/6000 with the *PowerPC* processor in 1990. With a clear focus on improving performance, reliability, and server utilization, IBM has continued to invest in all aspects of the POWER environment, developing new technology in the area of threads, virtualization, scalability, and availability to improve the utilization of server processing and enhance the concept of

[1] Complex Instruction Set Computer.
[2] Reduced Instruction Set Computer.

consolidation. (See Exhibit 1, at the right.)

No matter what your application mix - OLTP, BI, or HPC - IBM has what appears to be an unending string of industry-leading benchmark results. If we were to try to list all of the commodity or application benchmarks in which IBM was either #1 or in the top three, we would have to cut down a small forest, or at least one large Redwood. Keeping in mind that your application set will have unique performance characteristics, we will highlight the most significant commodity results for data center operation.

Most important for any enterprise with an OLTP application set is *transactional performance*. Here, IBM was so far ahead with POWER5, there did not appear to be anyone in second place. With a rating in excess of 3.2 million TPM, the 64-core *p5-595* had a commanding lead over the nearest competitor … the 32-core *p5-590* at 1.6 million TPM! The nearest non-POWER challenger was a 64-core HP *Superdome* at 1.2 million TPM. IBM also led in 2-, 4-, 8-, and 16-core core results, in price as well as price/performance. In terms of *high-performance computing*, IBM and POWER5 are again the clear leader using the *SPECompM2001* benchmark. The p5 holds the top positions at 4-, 8-, 16-, and 64-cores. If you are more interested in *application-specific* tests, IBM also led the *SAP SD 2-Tier R/3* benchmark with a rating of 23,456 users for a 64-core p5-595. Fujitsu is a close second here with a rating of 21,000, although that required a 128-core system. (Your results will vary.) This helps to explain IBM's recent success in migrating mission-critical database and business intelligence applications to AIX. Now, IBM has taken POWER5 and raised the ceiling.

POWER5+ Solutions

IBM has recently completed the introduction of a new line of POWER5+ models for System p. Having already announced dual- and quad-core systems at the entry and mid-range level, priced on top of x86 platforms but with more functionality, IBM has filled out the high-end of the p5 family with new SMP models for the *p5-590* and *p5-595* based upon POWER5+ for the largest data centers. POWER5+ provides even more scalability throughout the family, making the p5 an ideal consolidation platform for the SME as well as the largest enterprise. Using *rPerf*[3] as a measuring stick, a *p5-505* with a single processor has a rating of 4.10 while a *p5-505Q* with a *Quad-Core Module* (*QCM*) is at 20.25. At the high end, a p5-595 with 16 cores is rated at 108.13, while a 64-core model is at of 393.55, showing not only a healthy microprocessor architecture, but also outstanding linear scalability for virtualization and consolidation requirements.

Exhibit 1 – System p Technology

- **Advanced Virtualization** – Enables each CPU in the shared processor pool to be logically partitioned into as many as 10 micro-partitions and share Virtual I/O Server capabilities under AIX or Linux;
- **Multi-Threading** – Enables each core to run two application streams simultaneously, reducing the time to complete a task; this eliminates "dead cycles" by issuing two threads simultaneously and dynamically optimizing threads at runtime. Intel's *Montecito* has a single thread per core.
- **Mainframe RAS** – Delivers features normally found on larger, more expensive architectures including redundant service processor, Chipkill and ECC memory, First Failure Data Capture, dynamic de-allocation of selected resources, hot-plug slots, bays and power, and more.
- **Capacity-on-Demand** – Allows processors and memory to scale dynamically for temporary access for short-term needs or permanently to meet increases in long-term workloads.

IBM has now established a new ceiling for transactional performance, with an OLTP rating of over 4 million TPM, with price/performance of under $3/TPM, significantly better than any server so far based on Intel (*Itanium* or *Xeon*) or AMD *Opteron*[4]. These results show a dramatically better value proposition for any scale-up transactional application. This eclipses the previous leader – the p5-595 with POWER5 – by 25%, and nobody else is even close. To do this, IBM has taken advantage of R&D investments from across the POWER range, specifically from work on video game consoles, where highest performance graphics is a must. They have applied the same process – *Dual Stress Technology*[5] – to enterprise servers. This is a synergy unavailable to other server providers.

The p5 architecture, using a 90nm chip technology, is optimized throughout the system with faster processor-to-processor interconnects and faster processor to cache transfers to take advantage of advanced multithreading for a multi-core archi-

[3] *rPerf* is an IBM specific relative performance index showing the relative performance capability of any System p server as compared to a p610 with a *p3* processor.

[4] There are no results available for Sun's SPARC servers.

[5] Dual Stress Technology is a method of simultaneously stretching and compressing silicon to deliver a 24% transistor speed increase at the same power/heat level.

tecture, (i.e., 2 threads/core). The p5-595 is the perfect consolidation engine for the largest data center, with on-demand scalability from 16 to 64 POWER5+ cores running at up to 2.3 GHz, in increments of eight, providing up to 128 application threads to utilize its resources fully. With up to 254 micro-partitions, the p5-595 can access from 8GB to 1TB of DDR2 SDRAM memory, with on-demand scalability. The on-demand capability is available for trial, reserve[6], on-off, or permanent activations. In addition, many System p servers, such as the *p5-570*, p5-590, and p5-595 will be upgradeable to POWER6 in 2007 with commitments for twice the performance of POWER5, further energy efficiencies, a new release of AIX, and more RAS features.

In fact, IBM has made great strides in improving energy efficiency in the past five years with System p. In 2001, a p610 had a SPECint_rate2000 value of 7.7 at a maximum rate of .45 kilowatts, for a performance/kilowatt ratio of 17.1. In 2003, the pSeries 615 tested out at 16.9 for SPECint_rate2000, using the same level of energy for a ratio of 37.5. Today, the System p5-520 tests at 44.6 units with a maximum power consumption of .75, for a ratio of 59.4, significantly more than three times the energy efficiency.

In order to supplement the depleted IT staffs in many enterprises, IBM has implemented the *IBM Server Consolidation Factory for System p* to deliver a complete solution while lowering TCO. This enables data centers using a large number of smaller, less efficient servers to consolidate their applications, simplify management, and save on floor space and electricity. IBM can provide the solution framing, planning, design, and implementation that data center IT staffs can no longer provide. IBM *Tivoli Usage and Accounting Manager*[7] is available to measure, analyze, and report on the utilization and costs of a variety of computing resources from CPU to storage in a virtualized environment. This is an ideal application to manage charge backs to the appropriate department in a consolidated environment. IBM *Power Executive* is another cross platform tool used to monitor power consumption and assist in data center planning. All of these tools reduce overhead and improve the bottom line for any enterprise.

IBM has also established new *Value Unit* licensing for middleware. This includes an across-the-board 50% entitlement reduction for all System p Quad Core Module servers and significant savings in other models in order to achieve parity pricing with commodity x86 servers based upon Intel or AMD microprocessor technology.

Conclusion

Fortunately for the enterprise community, there is no end in sight for IBM innovation in improving performance and functionality, and lowering TCO, in the data center. The newest members of the p5 family are extremely powerful, scalable, and available systems, designed to enable a restructuring of IT to support a new business model and provide a roadmap for cost reduction. You may want to look at re-hosting your applications on Linux, on either POWER or an x86 platform, in the future, but for now, you can achieve a bigger bang for your buck with higher performance at a lower cost by retaining your existing operating environment

With unique virtualization capabilities and accounting features derived from years of mainframe experience, IBM has extended the consolidation capabilities of the highest performing open systems, transactional engine available. POWER5+ continues a legacy of price/performance leadership and is the leading open microprocessor, years ahead of Intel's *Itanium 2* (*Montecito)* processor. AIX5L features, such as dynamic logical partitioning, virtualization, and on-demand scalability, provide the flexibility you need to implement a consolidated data center. Whether you are one of many on the enterprise IT staff, or the only IT person in an SME, IBM has an *Express* solution to facilitate the redeployment of your IT infrastructure. They have enabled a simplification of the entire data center, not only through superior platforms, but also through outstanding programs such as the IBM Consolidation Factory, staffed with experts to improve your server utilization and reduce costs.

IBM is not only staying ahead of the competition, but they are playing leapfrog with themselves, constantly improving their own state-of-the-art architecture. If you already use POWER and AIX and need to improve performance and reduce cost, or if you use a competing solution that is not performing up to your expectations, take a close look at the p5 family – IBM could make you look like a hero.

[6] Reserve activation is only available for processors.

[7] See **The Clipper Group Navigator** dated July 3, 2006, entitled *Accounting as the Basis of Civilized Computing – IBM's Usage and Accounting Manager*, available at http://www.clipper.com/research/TCG2006053.pdf

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**, 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.