

IBM Raises the Ceiling for Application Server Performance

Analyst: David Reine

Management Summary

Taking a drive down the information highway can be quite different than taking a cruise along Interstate 10 in terms of capacity, performance, and cost of the ride, depending upon how many people (or clients) need to go along and how many vehicles (or servers) are available. If you are by yourself, an efficient, little moped scooter may be sufficient to get you from point A to point B. If you have a date, however, the back of a moped may not be the most glamorous position if you hope to get a second date. A little red convertible might be just the ticket. It is a bit bigger, a lot more expensive, a lot more performant, but does not have nearly the same efficiency. As the capacity requirement for transportation increases, say to transport a family of four, the capacity requirements change accordingly. Now the upgrade to a family sedan might be in order - not as sexy as the convertible, but certainly a higher capacity and, hopefully, more efficiency to fit into a family budget. As the transportation needs increase, i.e., a basketball team, perhaps, capacity requirements rise again. This time you may need multiple vehicles or one larger vehicle, a mini-van for example, to meet the transportation goal. Costs may go up and efficiency may go down. If you have to drive the Pop Warner football team, you may need anywhere from six to ten vehicles, or perhaps a bus, to get the job done. As the requirement for more passengers, or more clients, keeps rising, the size of the solution keeps going up, along with the price. Unfortunately, performance and efficiency don't always keep pace.

On the information highway, however, the story is quite different. As the number of application users increases, a simple upgrade to the current server will often suffice. More importantly, the performance of the platform and its efficiency can also improve, instead of suffering through the burden of additional capacity, depending upon the solution path selected.

If your data center chose the IBM *POWER5* architecture, for example, the Director of Information Technology (IT) can retain the original platform and upgrade in place to meet the ever-changing technology demands placed upon the IT staff, while improving the performance and efficiency of the platform, at the same time. With their latest announcement to the *eServer iSeries and pSeries*, the *i5-595*, *p5-590* and the *p5-595*, IBM has empowered the data center staff to replace an entire fleet of servers with a single platform, decreasing the total cost of ownership (TCO) and simplifying data center operation. To learn more about the additions to the *iSeries* and *pSeries*, please read on.

IN THIS ISSUE

➤ POWER5 Family Tree	2
➤ POWER5 High-End Configurability	3
➤ POWER5 Performance	3
➤ Conclusion	4

POWER5 Family Tree

It has been a very busy five months for IBM. In May 2004, IBM introduced the latest member of the *POWER* family, the 64-bit *POWER5* microprocessor, with their announcement of the *eServer i5*. Taking off from the position of strength where the *POWER4+* had positioned both the *pSeries* and the *iSeries*, *POWER5* added improved performance, improved capacity, and introduced to the *AS/400* community a level of virtualization that had previously only been available in a mainframe. In June, IBM presented the entry to medium-range members of both the *i5* and *p5* Series, ranging from mono- and dual-processor systems, to quad-processors, 8-way systems, and 16-way servers which stretched the previous limits of both commodity benchmarks and mission-critical application performance.

Designed to comply with the on-demand requirements of a new generation of IT workers, both the *iSeries*¹ and the *pSeries*² were introduced with minimalist configurations and maximum performance. They were positioned for expansion via a scalable platform design. However, the *p5* models introduced last summer - the *p5-520*, *p5-550*, and *p5-570* - already had gained leadership status in over 30 commodity benchmarks, including both high-performance computing and transactional processing categories. In fact, the 16-way *p5-570* outperformed a 72-way Sun *SPARC* server in an SAP benchmark.

In addition, *POWER5* had introduced virtualization features, such as micro partitioning, into both families for sharing a single processor between multiple applications and for creating multiple threads within a processor to, in fact, almost double the performance capability. These features enabled even small businesses with the flexibility to run multiple operating systems,

including *AIX 5L*, *Linux*, and *i5/OS*, simultaneously, within a *POWER5* server, virtually eliminating migration concerns for customers running with *AS/400* or *Linux* applications. In addition, the *i5-595* can run *Windows* and *Linux* on *Intel* through integrated *xSeries Adapters* and integrated *xSeries Servers*.

This new and innovative technology gave IBM's customers the performance, reliability, and scalability levels previously unattainable in affordable open systems server configurations. However, more importantly, **it did not disrupt the everyday, mission-critical capability of the data center; it merely expanded it.**

During September and early October, IBM continued to concentrate its marketing efforts at the entry levels, not only with server configurations, but also with storage. IBM has aggressively attacked the open systems *Linux* market with the *OpenPower 720*, an entry-level server aimed directly at their entry SMB customers who needed the virtualization capabilities of *POWER5* to reduce their operating costs and consolidate their processing requirements. IBM complemented that announcement with Express configurations for dual- and quad-processor models of both the *iSeries* and the *pSeries*, making these models even more affordable for SMBs.³ They also introduced a new family of low cost disk devices under the *TotalStorage DS* brand, the *DS300* and the *DS400*.

Just recently, however, IBM has turned its power and its *POWER5* back to the enterprise data center with the announcement of the *POWER5*-based *TotalStorage DS8000* storage subsystem⁴. Now, IBM is keeping the focus on the data center with the introduction of a new high-end for *i5*, the *i5-595*, and a pair of high-end *POWER5*

¹ See **The Clipper Group Navigator** dated May 4, 2004, entitled *IBM Introduces eServer i5 – A Recipe for Success* at <http://www.clipper.com/research/TCG2004039.pdf>.

² See **The Clipper Group Navigator** dated July 22, 2004, entitled *Open Systems and Virtualization – IBM Takes Another Stride with POWER5* at <http://www.clipper.com/research/TCG2004063.pdf>.

³ See **The Clipper Group Navigator** dated October 27, 2004, entitled *The Odd Couple Gets Engaged – IBM Walks the Walk with Linux for SMBs* at <http://www.clipper.com/research/TCG20040xx.pdf>.

⁴ See **The Clipper Group Navigator** dated October 18, 2004, entitled *True Enterprise Storage for On-Demand Data Centers – IBM Introduces TotalStorage DS8000* at <http://www.clipper.com/research/TCG2004086.pdf>.

application servers for the p5, the *p5-590*, and *p5-595*. Let us take a look at the capabilities of these enterprise class open systems servers.

POWER5 High-End Configurability

The business challenges for the p590 and both the i595 and p595 are the same. They must satisfy the identical business issues that face every other enterprise level server, simplify IT infrastructure, simplify management, and align the environment to respond to business priorities, except at a higher scale. Additional challenges include lowering the total cost of operation (TCO) while, at the same time, improving the reliability, availability, and serviceability (RAS) of the system. POWER5 provides an ideal base from which to meet these requirements on the grandest scales.

With the virtualization capabilities of POWER5 and the proven RAS characteristics of the pSeries, IBM has delivered an ideal platform for the consolidation of an entire data center onto a single platform. With micro-partitioning and the on-demand capability to transition system resources from one partition to another, the data center can move applications running under AIX, Linux, and i5/OS onto separate partitions of a single platform, simplifying the data center operation and lowering the TCO of the data center.

The p590 can be configured with as few as eight POWER5 processors running at 1.65GHz on a single multi-chip module (MCM), up to 32 processors with as many as half idle, available to be activated for on-demand processing. Starting with 8GB of memory, a single p590 can expand up to 1TB of DDR1 memory at 266 MHz or up to 128GB of DDR2 memory running at 533 MHz. Up to half of the memory may also be idle, in an on-demand architecture, waiting to be activated as needed in 1GB increments

Similarly, the p595 may be configured with as few as 16 processors and 8GB of memory, with expansion up to 64 POWER5 processors, up to 254 partitions, and up to

2TB of DDR1 memory or 256GB of DDR2 memory. The p590 can be expanded up to eight I/O drawers with 160 adaptors, while the p595 can support up to 12 drawers and 240 I/O adaptors.

Both systems come with a full set of virtualization tools, including Dynamic LPAR, micro-partitioning, shared processor pool, virtual LAN, and virtual I/O. The i5-595 can also expand up to 64 processors and 2TB of memory and the data center can upgrade the following models to an *i595*: *i825*, *i830*, *i840*, *i870*, and *i890*. This enables the data center to take advantage of the tens-of-thousands of ISV applications that are available under virtualization with the i5/OS operating environment.

POWER5 Performance

The very first thing that any CIO needs to understand is that the IBM *pSeries 690* with 32 POWER4+ microprocessors running at 1.9GHz, measures at over 1 million TPM in transactional performance benchmarks. *This is the #1 transactional mark for standalone systems, with greater performance than the HP Superdome with 64 Itanium-2 processors.* The Superdome is the only other standalone system measured at over 1 million TPM. If you include clustered systems, the p690 is ranked second against an HP rx5670, with 64 Itanium processors and 80 dual-processor Xeon clients. *With the p5-595, IBM is merely competing with itself to see how high they can push the performance bar.*

We do know the relative performance positioning within the pSeries for these servers, however. The p690 with 32 P4+ 1.9GHz processors has an IBM *rperf* value of 104.17. A 32-processor p590 running at 1.65GHz under AIX 5.3L has an *rperf* of 151.72. A p595 with 32 processors running at 1.9GHz has a value of 170.11. This is over 63% more performance. **The *rperf* for a 64-processor p595 with AIX 5.3L is 306.21; almost three times the performance of the p690.**

In order to measure adequately the relative performance gains of the i5-595,

IBM developed the Virtualization Grand Slam Benchmark to demonstrate the capabilities of virtualization for server and application consolidation. This test measures the performance of four applications running under different operating systems, i5/OS, AIX 5L, Linux, and Windows. IBM has shown that the i5-595 has over four times the performance of its predecessor, the eServer iSeries 890 with 32 processors.

We do have additional commodity performance data for the p595, from the **SAP SD 2-Tier Benchmark**. A 64-processor p5-595 has been measured at 20,000 users and 313 users/CPU - the #1 rating for users by far, and the #4 ranking for users/CPU, behind only three different configurations of the p5-570⁵. The nearest competitive system is the 128-processor Fujitsu *PrimePower* with a rating of 13,000, delivering only 102 users/CPU. This gives IBM a 53% edge over the nearest competitor.

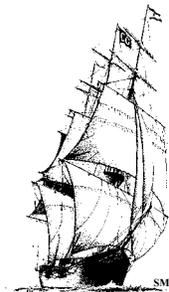
We also have a sampling of SPEC benchmarks to review. The first, the **SPECjbb2000 (Java Business Benchmark)** is SPEC's first benchmark for evaluating the performance of server-side Java. SPECjbb2000 continues the SPEC tradition of giving Java users the most objective and representative benchmark for measuring a system's ability to run Java applications. The 64-CPU p5-595 scores the highest on this benchmark at 2,200,162. A distant second is the 112-processor Fujitsu rated at 1,420,177, 54% lower than the p5. A second SPEC benchmark, **SPEC OMP (OpenMP Benchmark Suite)** is SPEC's first benchmark suite for evaluating performance based on OpenMP applications. SPEC OMP continues the SPEC tradition of **giving HPC users the most objective and representative benchmark suite for measuring the performance of SMP systems**. Once again, the 64-processor p5-595 scores highest on this benchmark with a result of 92,979. The

Fujitsu PrimePower HPC2500 comes in second with a result of 70,329 for 128 processors, 25% lower than the p5.

Conclusion

So, how should you look at these new POWER5 servers? Are they simply an upgrade to the current POWER4 environment, taking advantage of faster microprocessor settings while retaining binary compatibility with previous applications and enjoying the same RAS characteristics as older models? Alternatively, are they a major improvement in microprocessor architecture through the implementation of virtualization features previously not available? In fact, they are both.

IBM has presented a new architecture and two separate server families to take advantage of it. Whether the data center is primarily AIX (UNIX) or i5/OS (OS/400), IBM has developed an environment that is ideal for the integration of Linux and Windows applications within the framework of a dedicated, mission-critical application server. The 595, whether within the iSeries or the pSeries, can consolidate a data center into a single platform, improve overall application performance, and lower the total cost of operation over what you are doing today. Take a look to see how POWER5 servers might bring you closer to your goals.



⁵ These have an outstanding rating of 5,056 users for the 16-processor system and a #1 ranking in users/CPU at 328 for the quad-processor configuration.

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. 23. (Please dial "1-23" 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, 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.