

IBM Raises pSeries Performance Ceiling and Lowers Pricing Floor

Analyst: David Reine

Management Summary

The standard environment in the data center of 2004 is one of change. With corporate mergers leading the news on an almost weekly basis, we see data center managers faced with an ever increasing user base to attach to existing applications on installed servers, trying to hold the economic line on expansion. We see an era of simplification, with the data center manager trying to consolidate applications onto fewer servers in order to reduce the total cost of ownership. **Combine the two, expanding user count and consolidation, and processing requirements can quickly outstrip the computational and transaction processing capability of the average mission-critical application server.**

Furthermore, there is nothing more aggravating for the Information Technology (IT) Director than having to replace a mission-critical server because it has run out of horsepower. Typically, he has no problem with adding CPUs, adding memory, even upgrading installed CPUs with faster versions, **as long as the data center can retain its initial investment in the asset.** With a variety of servers to choose from, the IT Director needs to ensure that the architecture that he is acquiring is still in a growth phase. Twenty-year-old technologies about to enter end-of-life do not instill confidence when growth is a requirement of the business process.

IT management is looking for a solution with dynamic growth capability. This means a capability to grow in terms of CPU performance, turning on additional processors or replacing existing CPUs with newer, faster versions. This means the ability to activate additional memory in concert with the growth of connected users. It also means reducing the overall cost per user as a function of connections. To do this, the data center needs a platform that is scalable, flexible, and responsive to customer requirements, with a roadmap for future growth.

To accomplish this, IBM introduced a family of scalable servers, the pSeries (previously the RS/6000), designed to respond to a changing customer environment with the highest level of microprocessor performance, in a mission critical atmosphere. They have refreshed this family every two to three years, improving the CPU, adding memory. At the high end of the pSeries is the p690, introduced in 2001 to enable mainframe qualities in an open systems world, with the capability to consolidate all of a corporation's AIX and Linux needs into a single platform. Please read on to see how IBM has modified the p690 to make it the foundation of an on-demand computing infrastructure.

IN THIS ISSUE

- **EServer pSeries Overview** 2
- **P690 Enhancements**..... 2
- **Conclusion** 3

eServer pSeries Overview

In order to see the full value of the IBM *pSeries 690*, we first need to take a quick look back at the history of the *RS/6000* family and the *PowerPC* microprocessor. Initially introduced in 1990 as one of five RISC architectures in the market, the *POWER* architecture was originally composed of a set of seven or nine chips depending upon the size of the data cache, either 32KB or 64KB. The single chip *PowerPC* microprocessor was introduced in 1993 as a result of a partnership among IBM, Motorola and Apple to provide their customers a better price-performance ratio through the use of multiprocessing systems.

Because the new chip had to satisfy the needs of all three participants, it had to have the ability to:

- Permit a broad range of implementations to include low-cost controllers as well as high-performance processors;
- Include multiprocessor capabilities; and
- Define a 64-bit architecture that had binary compatibility with the 32-bit applications that preceded it.

Most importantly, **the architecture had to have a long life to sustain the investment that corporations would make in the mission-critical servers that would result.** Time has proven the decision to invest in the *POWER* environment to be a wise one. We can see in Exhibit 1 that IBM developed the *RS/6000* servers into an upgradable series of

Exhibit 1 - RS/6000 Upgradability

<i>u</i> proc	# of proc	Model	MHz	Year
POWER	1	320	20	1990
POWER2	1	58H	55	1993
PPC601	1	250	66	1993
PPC604	1	E20	100	1995
PPC604e	2	H10	233	1996
RS64	12	S70	125	1997
RS64 II	12	S70	262	1998
RS64 III	24	S80	450	1999
RS64 IV	24	S80	600	2000

Exhibit 2 – pSeries Upgradability

<i>u</i> proc	# of proc	Model	MHz	Year	rPerf min	rPerf max
POWER3-II	4	640	375	2000	1.00	3.47
RS64 III	4	660	450	2001	0.93	3.55
RS64 IV	4	660	600	2001	1.26	4.57
RS64 IV	6	660	750	2002	1.91	8.23
POWER4	32	690	1100	2001	18.02	60.66
POWER4	32	690	1300	2001	21.20	71.44
POWER4+	32	690	1500	2003	24.18	81.95
POWER4+	32	690	1700	2003	27.11	92.19

servers to provide continuing growth for those committed to its architecture over a ten-year period.

Starting in 2000, IBM recast the *RS/6000* into the *pSeries* with renewed life and higher scalability, as enumerated in Exhibit 2. Starting with the *POWER3* iteration of the architecture, IBM developed a relative performance scale to identify the additional computational capability of each succeeding model. The mono-processor version of the *pSeries Model 640* running at 375 MHz was given an arbitrary relative performance (rperf) value of 1, with each succeeding model measured against it. **We can see that the 32-processor version of the p690 running POWER4 at 1300 MHz has 71.44 times the processing power of the p640, true scalability.**

However, it is a rule of thumb that “all programs will expand to fit into available memory”. This means that no matter how much memory that the server can deliver, the application is always looking for more. The same holds true for processing power. Fortunately, the *p690* has plenty in reserve.

P690 Enhancements

In May of 2003, we reviewed IBM’s introduction of the *POWER4+* microprocessor along with their *Capacity on Demand* functionality¹. At that time, IBM announced the availability of 1.5 and 1.7GHz *POWER4+* processors and the capability to install 512GB

¹ See **The Clipper Group Navigator** dated May 14, 2003, entitled *IBM Soups Up Its pSeries Servers and Offers Temporary, Extra Power on Demand* at <http://www.clipper.com/research/TCG2003020.pdf>.

of memory. With the 1.5GHz processor, the p690 had an rperf² of 81.95; with the 1.7GHz processors the rperf was 92.19. In a more transaction-oriented measurement, the pSeries 690 at 1.7GHz was rated at over 765K TPM, which was the highest rated OLTP system at that time. In this world of leapfrog technology, however, that number was surpassed. This brings us to February 2004 with IBM's latest upgrade to the top of the pSeries family and a rise in the performance ceiling.

IBM has now turned up the performance of the POWER4+ to 1.9GHZ, while at the same time doubling the amount of memory that the p690 can support, up to one terabyte. The one terabyte consists of up to 8- 128GB memory cards which run at 633MHz to help ensure optimum memory subsystem performance, when used in conjunction with the 1.9GHz processors. The CPU boards also contain the higher-speed 128MB Level 3 cache that furnish an efficient interface between the processors and memory. In addition, you can upgrade any of the 1.1 or 1.3 GHz POWER4 processors or the 1.5 or 1.7GHZ POWER4+ CPUs with the new architecture. Furthermore, IBM has made both processor and memory capacity-on-demand features available to the p690+ family in 4 CPU and 4GB increments.

Incorporating the new POWER4+ processor at 1.9GHz into the benchmark generates an rperf of 104.17, which is the number that IBM now publishes for the p690+. However, this value does not incorporate the added benefit of an additional 512GB of memory that is now available. In an attempt to remove some of the extreme effects that additional memory would have on the rperf result, these latest numbers were benchmarked with 512GB of memory, the same level as with the POWER4. If an additional 512GB of memory were included, the results for the 690+ would be as follows.

Exhibit 3 – rPerf Values

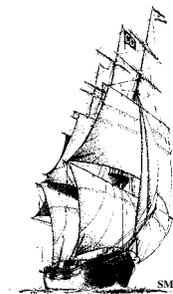
<u>uProc Speed</u>	<u>512GB Memory</u>	<u>1TB Memory</u>
1.7GHz	92.19	100
1.9GHz	104.17	113

This is a significant improvement in computational capability for the platform.

The question then becomes what is the relative improvement in the transactional performance of the pSeries 690+. We have seen that the p690 with 1.7GHz processors and 512GB of memory, using FAST drives, had a transactional performance of approximately 765K TPM. With 1.9GHz CPUs and 1TB of memory, that result improves to over 1,025K TPM, an improvement of 34% in OLTP, the highest single-system rating in the industry to date. Furthermore, the p690+ is economical as well, with a price/performance rating of just over \$5.40. **This is significantly lower than any other system in this category.**

Conclusion

With at least one eye firmly fixed on the horizon with *POWER5*, *POWER5+*, and *POWER6*, IBM continues to raise the bar on performance while holding the line on cost. With continued offerings in the area of Capacity-On-Demand, IBM demonstrates that they understand the advantages of a just-in-time delivery. **The p690 provides increased capability for both the high-performance computing applications reflected by the rperf performance ratios, and the high transaction counts and scalability required by OLTP applications.** IBM continues to deliver the kind of platforms that data centers require in the 21st century. Whether adding more users to current pSeries platforms or consolidating several servers into a single, new system, the extensions to the p690 deliver the performance required by either environment. If this sounds like your data center, take a look. The p690 may be perfect for you.



² Relative Performance (rPerf) is an estimate of calculated performance projections based on systems with maximum memory and the latest levels of AIX 5L and other pertinent software.

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.