



Right-Sizing the pSeries — IBM Tunes pSeries for Improved Configurations

Analyst: David Reine

Management Summary

Americans like things BIG. Moreover, we like them *faaaasssttt*. In the Halloween season, we hold contests to see who can grow the largest pumpkin. During the World Series, we watched to see who would hit the most home runs. Not singles, not doubles. We look for the BIG HIT.

When Americans go out for lunch – fast food – we make sure that we SUPERSIZE the FRIES. When we look to buy a car, we are concerned with how quickly it will go from 0 to 60 MPH. It does not even matter that most highways have a speed limit of 55 mph or less. Moreover, the car is not just a car; it is an SUV, a BIG SUV like the Cadillac *Escalade*, to claim the road – no matter the workload. What is the result of this excess? It is a waste of resources – in restaurants, it is a waste of food, or worse, a waste of health. On the highway, it is even a bigger waste of natural resources: a waste of gas and a loss of air purity.

What about the IT world? Normally the same story is true. We look for larger servers to provide more resources to applications to provide faster response times. We look for larger servers in order to consolidate slower and smaller platforms into a computing dynamo. If an eight-processor server is good, then a 16-CPU box must be better, and a 32-CPU server will be, well, SUPERSIZED. In addition, we continue to accelerate processor speed from 1.0GHz to 1.3GHz to 1.5GHz and beyond. Bigger. Faster.

However, in some cases bigger is not better. It is not necessary. In fact, it is inappropriate. In an era of on-demand functionality and just-in-time delivery, buying expensive resources to sit idle in a server where they will **never** be required is a waste of money on IT resources: CPU, memory, and storage. Installing an enterprise-sized platform in a department or branch is a classic example of over-provisioning, supersizing the server to excess, putting more resources in play than are required or, in fact, can ever be used. The cost of this over-provisioning is not trivial and should be of concern to the CFO of any size enterprise. How does the CIO find a platform that is properly sized for the business environment in which it is installed, and can still provide an efficient application platform with more than adequate response times? One thing he can do is to check out IBM's *pSeries* and the diminutive, but accelerated *Model 615*. To find out how IBM might help you find the platform sized to your environment at both the entry and enterprise level, please read on.

IN THIS ISSUE

➤ PSeries Server Family	2
➤ pSeries 615 – Reliability for the Entry Level	2
➤ On-Demand Functionality for the High-End	3
➤ Conclusion	4

pSeries Server Family

During the turbulent '60's of the 20th Century, the social theme of the day was "Power to the People". As we progress into the 21st Century, it would appear that IBM has modified that message for the distribution of their *pSeries* family of *AIX* and *Linux* servers. Now it is "**POWER4 (for) all of the People**".

Enterprises in the 21st century have offices and computing centers around the world. They consist of the "glass house" data center within the corporate headquarters, distributed processing centers in regional offices, and even computing nodes of varying size in branches located at the local level to service individual needs. While the applications residing on these servers need to be available on all platforms, regardless of size, the computing requirements in the local office and remote backup sites do not approach the provisioning requirements for the enterprise server in the data center. **The pSeries attempts to address the needs of each of these computing nodes.**

In order to provide the required performance level for the data center, IBM has introduced, and continues to enhance, the high end of the pSeries product set, the *pSeries 670 and 690*. In response to the increasing demands from IT for computing power, IBM has introduced platforms with ever-expanding CPU count, ever increasing CPU performance and reliability. Today, the pSeries 690 supports from eight up to (32) 64-bit *POWER4+* microprocessors running at speeds of up to 1.7GHz, with the capability of logical partitioning (LPAR) to subdivide that capacity among many applications. This open server platform has mainframe performance of over 750,000 TPS in support of thousands of users located over a wide geographic network and a relative performance (rperf)¹ factor within the pSeries of 92.19. Moreover, with an

¹ *rperf* is an estimate of commercial processing performance relative to other pSeries systems. The IBM pSeries 640 model B80 is the baseline reference system and has a value of 1.0

availability now measured at 99.999% for both the p670 and the p690, an enterprise can count on the reliability of the pSeries.

For regional offices, IBM has introduced the *pSeries 630 and 650*. With the ability to support two-to-eight *POWER4+* processors, the pSeries 650 has the same architectural characteristics of its big brothers, the p670 and p690. With an rperf value of 18.67 for an 8-way server, the p650 is perfectly positioned to handle the workload for regional offices with hundreds of users.

With the latest announcement, **IBM now addresses the remote branch office. With significantly lower concurrent demands, the branch can now be equipped with a server with the same copper and silicon-on-insulator *POWER4+* architecture as the enterprise and regional systems. Moreover, it can have the same *AIX* or *Linux* operating system.** The introduction of 1.45GHz processors for the *pSeries 615* raises the rperf for this entry-level server from 4.0 to 4.41 for a dual processor configuration, providing the p615 with mid-range performance at an entry-level price point. This performance, proportional to the pSeries 650, and to the p690², enables the CIO to maintain application compatibility across the enterprise while lowering the TCO of the branch systems to meet the constraints of 2003 budgets.

In addition, new configuration rules and pricing at the enterprise level enable the distribution of fully configured remote backup servers at a significantly reduced acquisition cost.

pSeries 615 – Reliability for the Entry Level

Configurable as either a mono- or a dual-processor system, the pSeries 615 can be targeted at distributed high-performance workloads in a small or medium business (SMB) environment. It includes many of the same reliability, availability and serviceability (RAS) features as larger pSeries

² On a per processor basis

**Exhibit 1 –
P615 Reliability Features**

- First Failure Data Capture
- DDR ECC Chipkill Memory
- Bit-Steering/Redundant Memory
- Memory Soft Scrubbing
- Redundant Power/Fans
- Dynamic Processor Deallocation
- ECC PCI Bus, L2/L3 cache
- Hot-plug PCI Slots, Fans, Power
- Internal Light Path Diagnostics
- Hot-swappable Disk Bays

servers. (See Exhibit 1 on the next page.) Already configurable with 1.2 GHz CPUs, **the addition of 1.45 GHz CPUs and new I/O make the p615 the ideal platform for mission-critical applications at the branch level.** The new I/O includes an interface for up to eight internal Ultra320 disk drives, at speeds of up to 320MB/sec. Ultra320 SCSI, a new industry-standard storage interface, transfers data over a 16-bit LVD (low voltage differential) bus at speeds up to 320 MBps.

Available in a deskside package, or rack-mountable in a 4U space, the p615, a 64-bit system, comes with 1GB of memory, but is capable of addressing 16GB of memory. **This is sufficient to keep most databases accessible in high-speed memory, enabling applications to run faster by limiting the number of physical disk access.** The memory subsystem consists of error checking and correcting (ECC) memory to detect both single and double-bit errors, and can correct all single-bit errors dynamically. *Chipkill bit-steering* memory complements the ECC functionality by automatically activating memory spares when it detects multiple memory errors. **This significantly reduces the number of memory failures, thus reducing the number of system outages and improving the overall**

reliability of the p615.

The *p615* comes with four disk devices standard and four optional, additional devices. These devices are available in a variety of capacities, 36.4, 73.4, and 146.8GB, all supported over a 320 MBps bus, and mounted in hot-swappable disk bays, for an internal capacity of over 1TB. This hot-swap-ability allows the replacement of failed drives without having to power the system down. RAID0, RAID5 and RAID10 options are available internally using the new PCI-X Dual Channel Ultra320 adapters, enabling improved performance and increased data reliability, less expensively than the previous adapters. **This relieves IT from the necessity to configure external disk devices in order to achieve this RAID reliability, reducing the TCO and improving the ROI.**

Additional RAS functionality is standard in the form of an integrated service processor that monitors system health. This feature detects error conditions and automatically places a service call to IBM, often before the condition is apparent to systems administrators or users. In fact, **the service processor can initiate a dynamic system reconfiguration to bypass the failure and keep the system operational.** This reconfiguration helps to eliminate outages and maintain continuous application availability. In addition, the service processor facilitates remote systems management, permitting remote shutdowns and reboots. It also provides support services for the installation of application upgrades and operating systems patches, again remotely.

What does this mean? **It means that the enterprise does not need to have a systems administrator at every branch or local office to provide mundane systems services.** A single administrator in the data center can manage the server, saving the enterprise money on the IT budget. What else does it mean? **It means that by combining mid-sized performance, an entry-level price, and enterprise resiliency, IBM has made the p615 the ideal server for the branch office.**

On-Demand Functionality for the High-End

Even with the attention paid to the improvements to the pSeries at the entry level, IBM has not forgotten the enterprise level. Although the base platforms have remained static, IBM has enhanced the programs surrounding them with a smaller initial configuration and improvements to on demand functionality. Specifically, IBM has improved the configurability of the *p670* to enable installation with only four active processors, **reducing the provisioning cost**. Before, a minimum configuration required eight active processors. In addition, IBM has introduced a new program for Capacity Backup for the *p670* and the *p690*, for disaster recovery. Moreover, *Linux* as well as *AIX* now supports Capacity Update on Demand for these models.

In order to make the installation of a pSeries 670 or 690 more attractive, IBM has instituted new configurability, and therefore new pricing, for both of these models. The data center may now enable, for example, a *p690* as a Capacity Backup server with a single active MCM (Multi-chip Module) containing four active and four inactive POWER4+ CPUs, and up to three additional inactive eight-processor MCMs. Previously, the requirement was for 8 active CPUs with 8 inactive, or 16 active CPUs with 16 inactive. The pSeries 670 requires 1.5GHz processors, while the *p690* may use the 1.3, 1.5, or 1.7GHz CPUs.

Because of this change in policy, the CFO can reduce the initial cost of a Capacity Backup server, significantly. For example, if IT chooses to install a pSeries 690 with 32 POWER4+ CPUs as a primary server, they can configure a 32 CPU Capacity Backup server with only four active CPUs for \$920K. The typical cost of a server with 32 active CPUs is \$1.95M, **thus deferring approximately 75% of the cost until resource activation is required by the failure of the primary server and the need for on-demand business continuance**. In addition, each additional MCM comes with 240 processor days of on/off activation.

This means that a fully configured *p690* comes with 720 processor days of on/off CPU activation. The permanent activation of CPUs on a Capacity Backup server is not available. A complete list of services that comes with the acquisition of a Capacity Backup server appears in Exhibit 2.

Exhibit 2 - CuOD Services

- Initial setup test period
- Annual test time
- Instant activation of dormant processors
- Ability to test dormant processors on a regular basis
- 240 CPU days of processor activation per MCM
- No monitoring or reporting required
- IGS services available on an as required basis

Conclusion

In an era where the costs of over-provisioning your application server can literally destroy your bottom line, IBM has addressed the issue head on. At the entry level with the *p615*, they have equipped a dual-processor system with the same CPU and I/O performance as the largest servers in the pSeries line. Significantly, **they have lowered the price/performance point for the remote branch or distributed department application server**.

At the high end, they have reduced the acquisition cost of disaster recovery, lowering the price for idle resources. **By introducing a reduced active configuration for the Capacity Backup server, IBM has enabled the CIO to protect his investments today, while deferring charges for that protection until it is required. IBM absorbs the cost of over-provisioning today, in order to proliferate the capability for tomorrow.** For most enterprises, this is a good deal.



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 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. 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. 37. (Please dial “1-37” 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.