THE CLIPPER GROUP
Navigator™
Navigating Information Technology Horizons℠
Published Since 1993 Report #TCG2004063 July 23, 2004

# Open Systems and Virtualization — IBM Takes Another Stride with POWER5

Analyst: David Reine

## Management Summary

Some enterprises are so savvy that they get it right the first time, every time. They are rare. Most feel fortunate if they get it right most of the time, just as long as they don't get it really wrong the rest of the time. Nobody wants to be the business that gets it wrong all of the time. They are not around for long.

Take the automotive industry for example. Ford is pretty good; they get it right a lot. Take the *Thunderbird*, for example. Debuting in 1954, Ford continues to sell the T-Bird with variations and new features, but for the past 50 years, it has stayed true to the spirit of the original. Why not, Ford has sold over four million of them. Another example of Ford getting it right the first time is the *Mustang*. Introduced in 1964, the Mustang has been "refreshed" annually, with more than a half-dozen generations. Production of the Mustang continues today. If it's not broke, don't fix it! Then there was the *Edsel*. Introduced in 1957 as a medium-priced family car, it was discontinued in 1957. Oh well, you cannot get them all right.

Another example of getting it right would be the *Volkswagen Beetle*. Originally introduced in 1938, VW has sold over 21.3 million "bugs", although representatives from Ford reviewing plant production in Germany after WWII said that it "was not worth a damn". Probably the same guy who introduced the Edsel. The Beetle did not remain the same over the years. Like the T-Bird and the Mustang, VW gave it new features and produced several variations. The Bug was discontinued in 1978, but like any classic, demand was so great that VW brought it back in 1998.

There are examples of getting it right in high-tech also. None is more notable than IBM's development of the POWER technology. In an industry where some technologies are outdated within 18 months of introduction, the legacy developed by IBM is worth noting. With its RISC roots dating back to 1986 with the RT Personal Computer, IBM evolved the technology into the RS/6000 in 1990, with a multi-chip *POWER* processor package known as *POWER1*. Over the past decade+, IBM has further evolved the architecture down to a single chip, transitioning from 32-bits to 64-bits along the way. *POWER4* was introduced in 2001 with two 64-bit *PowerPC* cores on a single die. Now IBM has improved the technology one more time with the *POWER5*, or *p5*. If it ain't broke, don't fix it. To learn more about the p5 and the newest members of IBM's "p" family, please read on.

## IN THIS ISSUE

- Data Center Requirements in 2004 ..... 2
- POWER5 Processor Innovation ..... 2
- eServer p5 500 Family ..... 4
- Conclusion ..... 5

## Data Center Requirements in 2004

In 2004, data centers have begun to take on a very disconcerting appearance. After making significant network-wide expenditures in the late 1990's in order to prepare for Y2K, a downturn in the economy has left many data centers with an antiquated environment and an aging, but still functional, fleet of application servers. In order to protect the enterprise and ensure that the mission-critical applications would run reliably, the staff has isolated the servers as small islands. With a single application running on each, the data center has seen the proliferation of poorly utilized application servers within the enterprise. Still, the CPU performance capability of each of these servers exceeds the demand on these processors. This has resulted in a network of servers that are idle from 50-70% of the time. The costs associated with this inefficiency are staggering.

The data center must expend resources to manage every server in the network. IT has to manage data stored internally and externally, via direct connections, for each. The CIO faces mounting bills for maintenance for servers that have excess processing power, but no means to tap into it. Not to mention the cost of energy to power the systems and to cool the air, as well as the cost for the floor space the servers occupy.

The IT staff has considered upgrades to the servers to enable multiple applications to share the excess processing time. However, in many cases, the microprocessor architecture at the center of those servers does not contain the necessary capability to run different operating systems to handle the different application requirements. In other instances, the microprocessors are obsolete and not even available. Even if upgrades were available, does the data center want to consider throwing good money after bad chasing a terminal architecture? They have considered consolidation plans before. However, with the constant advancements made to processors and servers, no one was ready to put a stake in the ground in order to commit to a specific technology.

What are the characteristics required by today's on demand business generation for an application server? Some would consider better system utilization to be high on that list.

### Exhibit 1 – Characteristics of a Superior Server

- **Flexibility** – The ability to execute multiple operating systems, such as UNIX, Linux, Windows, or other important legacy operating systems.
- **Performance** – Recognition within the industry as a leader in one of the standard benchmarks, such as TPC, SPEC, or SAP, among others.
- **Reliability** – A track record with a high MTBF and low MTTR. Untested architectures need not apply.
- **Manageability** – The emphasis here is on ease of use – a single pane of glass to manage multiple environments.
- **Scalability** – The ability to upgrade, in place, in an on-demand manner, from mono- to multi-processor and from one generation to the next.
- **Affordability** – The cost of the system must be within the budget of the average SMB, not just the largest enterprise.

Outstanding system performance is certainly a requirement, as is reliability and manageability. Moreover, when you get to the bottom line, the number that you see there must be affordable. (See Exhibit 1, above.) The next server that the data center acquires must enable multiple workloads with flexibility and efficiency. The question then becomes, where can the CIO find a solution that meets these criteria?

IBM has recently announced a new generation of servers within their *eServer* family. These servers, the *p520, p550*, and *p570* are all based upon the *POWER5* microprocessor technology, the latest upgrade to the 15-year run for the *POWER* architecture. Before we examine the new server family, let's take a look at the microprocessor itself, to see how IBM has enhanced its core technology.

## POWER5 Processor Innovation

Unlike the strategies that some of their

competitors are implementing, IBM has chosen to introduce new, innovative microprocessor ideas while maintaining binary compatibility within their highly performant POWER technology. Designed to enable virtualization within the processor in order to improve the efficiency of UNIX/Linux servers, IBM's *p5* allows a mix of applications to share the same processor and resources through the IBM *Virtualization Engine*. IBM does not disrupt their customer's mission-critical operations by transitioning to a brand new architecture. Combined with *AIX*, POWER CPUs have delivered over the past decade the consistent performance required to keep IBM's *pSeries* servers in the forefront of IT development. So much so that IBM now shares the p5 platform with the *iSeries,* which just recently was announced as a POWER5-based *i5* server[1]. This continues a trend initiated 10 years ago by moving *AS/400* applications (now *i5/OS*) to the POWER family. Overall, based upon the benchmarks run, POWER5 appears to have twice the performance of Itanium-2 servers. Some of the more significant improvements made to the POWER5 architecture for virtualization, workload management, and reliability are:

- Simultaneous Multi-Threading (SMT) to make each CPU look like two to AIX 5L V5.3;
- Automatic dynamic engagement of inactive on-demand processors based upon pre-defined enterprise objectives;
- Enhancements to the *Hypervisor* to manage the micro-partitioning capability in Dynamic Logical Partitioning (see Exhibit 2);
- Cross-partition policy-based workload management;
- Virtual storage and I/O in a micro-partition environment
- Hot I/O drawer addition/removal; and
- Dynamic power management for CPUs.

Through Micro Partitioning, the p5 can run *AIX* 5L V5.3, *Linux*, and *i5/OS*[2] not only on the same system, but also on the same processor, at the same time. In fact, the data center can execute multiple iterations of AIX, enabling the execution of production software at the same time that development/test continues on later versions. AIX 5L V5.2 does not support Micro-Partitioning, SMT, or virtual I/O.

Sharing the basic infrastructure between the p5 and the i5 also should enable earlier access to the integrated middleware. More third parties will take advantage of the synergy of the common POWER5 hardware to port their applications to an IBM environment.

**Exhibit 2 – Micro-Partitioning Functionality**

Derived, in part, from the zSeries, IBM's *Hypervisor* is the glue that holds the components of the *Advanced POWER Virtualization* together. Advanced POWER Virtualization is an option for p5 systems including the firmware enablement for micro partitioning and the *Partition Load Manager*. The Hypervisor allows the dynamic partitioning of processing resources down to $1/10^{th}$ of a processor, with increments of $1/100^{th}$, and prevents partitions from accessing resources or data owned by other partitions. It also enables the dynamic assignment of idle processors to any application running any of the qualified operating systems in an on-demand fashion. However, because LPARs (logical partitions) sit on top of an SMP architecture, some resources, such as I/O, are also implicitly shared by the partitions. The Hypervisor keeps the partitions from using shared resources in a way that would deny or restrict access to those resources by other partitions. Micro-Partitions still require dedicated memory.

## eServer p5 500 Family

IBM recently has introduced four models into the eServer p5 lineup, two at the entry

[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] i5/OS will be available on p5 servers in 1H05 and is limited to one processor on the p5-570 and up to two processors on select high-end models.

level and two in the mid-range. At the entry level is the *p5-520* dual-processor and the *p5-550* quad, aimed at SMB and enterprise departmental requirements. Either of these platforms is ideal for the organization that wants to consolidate a small, diverse workload onto a single processor. In the mid-range are the 8-way p5-570 Express and the 16-way p5-570 aimed at enterprises that plan to combine virtualization with consolidation, in order to simplify the overall enterprise IT structure. The IT staff can configure these models to take advantage of IBM's capacity-on-demand (*CoD*) capability to allow the enterprise to respond rapidly to changing needs. IBM plans to make an announcement at the high-end of the family later this year with the expectation that this will allow them to continue two years of consist growth into a third. AIX 5L V5.3 is designed to support 64-way SMP configurations (up to 128-way with SMT enabled).

All of these systems can use the IBM Hypervisor to create the flexibility that the servers needs to employ better utilization of the processing capabilities through partitioning and dynamic resource movement. They can also support virtual I/O to provide the ability to dedicate adapters and devices to a virtual server rather than a physical one, allowing the on-demand allocation of I/O devices. These virtualization features are available as an option for $590/CPU on a p520 or 550 and $990/CPU on the p570.

The performance of this new family of servers matches outstanding features with impressive improvements in performance. Using IBM's relative performance (rPerf) scale for measuring POWER servers, we can see (Exhibit 3) that IBM has leap-frogged their previous performance levels. With a 125% improvement at the entry level, from 4.41 on a p630 running at 1.45 GHz to 9.86 for a p5-570 running at 1.65 GHz, IBM has more than doubled the performance of an entry server. Moreover, there is a 150% improvement compared to a p5-570 running at 1.9 GHz.

At the top of the mid-range, with 16 processors, we see an extremely significant rPerf bump from a p670 Turbo to p5 with a 65% improvement to 77.45. These are a result of the SMT feature of AIX L5 with V5.3.

Comparing the p5 set with competitive systems is equally impressive. Using the SPEC floating-point benchmark as a basis, we see results in Exhibit 4, below, that are even more remarkable.

**Exhibit 3 –**
**Relative Performance (rPerf) Index**

| | Dual | Quad | 8-way | 16-way |
|---|---|---|---|---|
| p630@ 1.45GHz | 4.41 | | | |
| p670@ 1.5GHz | | 13.66 | 24.18 | 46.79 |
| p5-570 @1.65GHz | 9.86 | 19.66 | 37.22 | 68.40 |
| p5-570 @1.9GHz | 11.16 | 22.26 | 42.14 | 77.45 |

In the single CPU SPEC *fp2000* benchmark, a mono-processor p570 outperforms an HP *rx4640* with Itanium-2 by 25% with *Opteron* and *UltraSPARC III (Sun Fire V480*) trailing badly. Looking at the multi-processor *fp_rate2000* benchmark, we see the results for 16-way servers. The p570 leads the parade with an outstanding rating of 460. A 16-processor HP *Superdome* is next with a rating of 235, just about ½ the value of the p5 system. In fact, the data center would need a 32-way Superdome with a rating of 470 to best the p570. But, at what price?

Results from SPEC's *Open MP* benchmark, measuring system performance for real scientific and engineering applications, showed similar results with a 16-way p570 rated at 38,282 vs. 19,281 for a 16-way Superdome. With a 32-way Superdome measured at 29,106, it would take a 64-way system at 42,290 to outperform the p5.

**Exhibit 4 –**
**SPEC Floating Point Results**

| | GHz | fp2000 | fp_rate2000 |
|---|---|---|---|
| *Power5* | 1.9 | 2702 | 460 |
| *Itanium-2* | 1.5 | 2161 | 235 |
| *Opteron* | 2.4 | 1553 | |
| *uSPARC* | 1.2 | 1084 | 174 |

### *eServer p5 520*

With POWER5 processors running at 1.65GHz, the p5-520 is configurable as a mono- or dual-processor, in a rack, or desk side, and can support up to 32GB of memory and 32MB of L3 cache. The p520 has achieved leadership position on the SPEC *OpenMP2001* benchmark with almost twice the performance of the HP *rx2600*.

### *eServer p5 550*

With twice the configuration of the p520, the p550 supports four 1.65GHz p5 CPUs and up to 64GB of memory, in a CoD mode. The p550 also supports 72MB of L3 cache.

### *eServer p5 570 Express*

The p5-570 Express is a rack-mounted server with support for up to eight POWER5 processors at 1.5GHz and up to 256GB of memory and 72MB of L3 cache. As an *Express* version, this model is sold on the web for the SMB or enterprise account, and there are discounts for purchasing standard configurations. On-demand features are not available on the p570 Express.

### *eServer p5 570*

The p5-570 server, also rack-mounted, supports up to 16 POWER5 processors at 1.65 or 1.9 GHz and up to 512GB of memory. The p570 is a true CoD server with IBM enabling the delivery of four active and up to 12 inactive CPUs, with inactive processors and memory capable of activation. This is where IBM really shines in performance testing.

**In transactional performance rankings, the IBM p5-570 is the only 16-way server among the top 10 systems tested**. The rest are either 32- or 64-way *Itanium-2*, 64-way *SPARC64V*, or 32-way POWER4 servers. POWER5 not only leapfrogs POWER4, it also dwarfs the competition. The current 16-way p5 processor ranks third in performance, but it ranks first in price/performance at $4.95 per tpmC[3]. Based upon current projections, the proposed 32-way p5 should provide 50-60% more performance than today's 64-way Itanium boxes. Using the SAP *SD 2-Tier* benchmark, a 16-way p570 has the same power as a 36-way (72 cores) *E20K* from Sun.

Because IBM has not yet announced a high-end server, they will keep the pSeries 670 and 690 servers alive through 2004. In fact, IBM has implemented a plan where it will be cheaper to buy a p670 or 690 now and upgrade it to a new p5 server later, than waiting to buy the new server in 4Q04.

## Conclusion

IBM tosses a new generation of UNIX/ Linux servers into the consolidation ring with p5. This engine has remarkable performance and unmatched flexibility in terms of operating system compatibility and scalability. The p5 family has advanced virtualization features, unavailable with competitive servers, to free the IT staff from repetitive administrative functions. Micro-Partitioning allows for the increased utilization of system resources by automatically applying only the required amount of processor resource needed by each partition. Moreover, Tivoli agents may be employed to provide comprehensive network security and intrusion detection.

The p5 has on-demand capabilities that others do not even offer. IBM offers binary compatibility for your mission-critical applications while others offer training for new platforms when they obsolete your enterprise server. IBM will also help you transition from your old system to a POWER server. Through Global Services, IBM offers a free migration toolkit to enable most applications to move to POWER. With POWER5, you will have clear sailing with the wind to your back. IBM can empower you to get it right, this time.

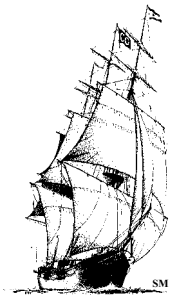

[3] Based upon figures provided by IBM as of July12, 2004,, an 8-way p570 measures 371K TPM and a 16-way measures at 809K.

### 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



### 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.