sageza

Instant Insight
May 21, 2007

IBM Announces POWER6 and Updated System p 570

By Clay Ryder

IBM today announced its latest Power architecture-based processor, the POWER6.

At 4.7 GHz, the dual-core POWER6 processor doubles the speed of its predecessor, the POWER5, while using approximately the same amount of electricity. The chip will be first delivered in a new version of the IBM System p 570, that claims the leading position in four widely used performance benchmarks for UNIX servers: SPECint2006 (integer calculations), SPECfp2006 (floating point calculations), SPECjbb2005 (Java performance in business operations), and TPC-C (transaction processing), an industry first. Additionally, the new system holds twenty-five benchmark records across a broad portfolio of business and technical applications.

Key features of the new IBM System p 570 with POWER6 include:

- Dual-core POWER6 processor in three speed ratings of 3.5Ghz, 4.2GHz, and 4.7Ghz
- 8MB of cache per POWER6 chip
- Hardware-based decimal floating-point arithmetic
- Live Partition Mobility
- Live Application Mobility
- x86 architecture Linux applications support
- Enhanced energy management

The POWER6 processor clocks can be dynamically turned off when there is no useful workload and restarted when there are instructions to be executed. Power saving can also be achieved through dynamically powering off and on memory when needed. In cases where an over-temperature condition is detected, the POWER6 chip can reduce the rate of instruction execution to remain within an acceptable, user-defined temperature envelope. Additionally, the separation of circuits that cannot support low-voltage operation onto their own power supply rails allows significant power reductions for the rest of the chip depending upon workload or performance needs at a given moment.

Net/Net

In an era where raw technology itself is less important than the business applicability of a IT solution, it would be easy to tend to overlook the latest version of a given processor. After all, are not all processors more powerful than their predecessors and accompanied by an ever-decreasing cost/performance curve? In general, this is true; however, it is also a gross simplification. While commodity processors from Intel today boast incredible performance, RISC and other high-performance processors continue to push the performance envelope and innovate in ways that are not often found in commodity offerings. The new POWER6 processor is a case in point.

Some of the top issues IT and datacenter managers are facing today include underutilized resources, power and cooling limitations, bloated operations budgets, and growing demands for more applications and data

The Sageza Group, Inc.
32108 Alvarado Blvd #354
Union City, CA 94587
510·675·0700 fax 650·649·2302

storage. Consolidation of servers through virtualization schemes has been posited as a solution to the cost-ineffective infrastructure found in many organizations today. While this is a good start towards solving the inefficiencies of IT infrastructure run amok, the reality for most mid-sized or larger organizations is that the virtualization solutions afforded by software-based endeavors can only go so far. This is in part due to the historic focus of commodity processors on ever-increasing clock speeds, not the holistic nature of the systems in which they are incorporated. Contrast this with IBM's assertion that its performance leadership is due to the balanced design that scales POWER6's performance along with system design criteria such as cache sizes and bandwidth in a balanced way. This is where the value of a holistic approach to system design as evidenced by the System p 570 becomes evident.

POWER6 was designed with virtualization and consolidation as a driving consideration. The hardware-based hypervisor affords efficiencies and separation of operating system components that are not possible in the more common software-based virtualization schemes. While software-based approaches have considerable merit, they tend to not scale to the degree necessary for larger if not massive consolidation initiatives. When one considers the many common workloads that are ripe for consolidation such as Web-oriented infrastructures based upon Linux and open source running on largely uncustomized x86 servers, the need for a highly scalable virtualized server environment starts to come to life. Through the recently announced System p AVE emulation, many of these common binaries can be directly consolidated onto a System p virtual environment with no mandated recompile. Given the sheer compute and throughput power of the p 570 with POWER6, for many organizations the consolidation of all of their x86 systems could take place on a single System p. The implications for reduced operational expense are considerable.

The energy efficiency angle of the new chip and server are also reflective of a more holistic system design. The inherent ability of POWER6 to throttle back power consumption in certain scenarios and shut it down completely in others is leveraged by software such as PowerExecutive to provide a very granular control over energy consumption within the System p. This extends beyond the helpful, but relatively straightforward, notion of shutting off memory or cores that are not executing a workload at a given moment. Rather the ability to slightly reduce performance by 15% but reduce power consumption by 50% implies a whole new set of usage scenarios where energy consumption can become a more significant factor in determining workload placement. In less than fully loaded environments, the performance of POWER6 compared with POWER5 more than offsets the slight performance reduction incurred with significantly reduced power consumption. Thus, the POWER6 could effectively offer twice the performance of POWER5, not at half the power consumption, but at closer to a third or less.

This ties into one of the very compelling aspects of the new p 570, namely Live Partition Mobility. This ability to move virtual partitions seamlessly—i.e., without suspending them followed by a reboot—between any suitably equipped System p allows administrators to treat servers as a pooled resource as opposed to a workload specific model. In addition, partition mobility bolsters the energy efficiency discussion even further. Just as electric utilities constantly monitor generation facilities to optimize output per minimum production cost, so too will IT personnel be able to move workloads to the systems that can bear the workload at the minimal cost per computational unit. Being able to adjust workloads to maximize performance while minimizing power consumption brings a new factor into datacenter planning. Clever deployment and redeployment of workloads could result not only in notable energy savings, but also in averted capital expenditures to augment data center floor space, cooling, or electrical capacity.

Overall, we believe the POWER6 is more than an incremental improvement on a venerable architecture. The potential to do more with less is well illustrated, but so too is the potential to revisit the notion of physical as well as environmental/energy footprints. In addition, the possibility of consolidating many x86 Linux binaries onto a single physical server supporting numerous virtual Linux environments proffers improved operational/management efficiency. To our way of thinking, POWER6 is more than just a new and improved processor; it illustrates strategic value of R&D investments combined with an out-of-the-box approach to solving present and future IT challenges.

The Sageza Group, Inc.
32108 Alvarado Blvd #354
Union City, CA 96587
510·675·0700 fax 650·649·2302