IBM's Power: A View from the Hill

By Gordon Haff

Introduction

With a pair of April announcements, IBM has largely completed a major transition of its Power Systems lineup. Its POWER6 processor is now available from top to bottom. However, these announcements, together with a variety of others over the past year, weren't merely about new chips, new servers, and new speeds and feeds. They also concerned sweeping organizational changes within IBM's Server and Technology Group (STG); the merger of System i and System p into a single Power server line; a significant broadening of IBM's virtualization portfolio; and much rationalizing and revamping of names, logos, and the like.

In short, we've seen the most significant changes to the way IBM STG runs its business and presents itself to customers since 2000, when it corralled its historically fragmented server lines under a single organization and R&D structure, publically identified as the umbrella "eServer" brand. So what better time than the present to climb up to a good vantage point and take a close look at IBM's Power landscape—the servers themselves, of course, but also the processor technology, systems software, and virtualization capabilities that complement them?

Drawing the Map

Before pulling out the binoculars to take a closer look at the details, it's useful to survey the changes more broadly.

Overall, the Power Systems hardware line is more streamlined than the previous generation. For example, all servers run standard POWER6 processor modules, whereas the POWER5 generation included Quad-Core Module (QCM) variants offering more cores running at a lower frequency than the standard Dual-Chip Module, in addition to using a Multi-chip Module (MCM) at the high-end. And the Power 595 now anchors the high end of the line by itself, where previously there was both a p5 590 and a p5 595. IBM has also not announced POWER6 replacements for two entry-level rackmount servers; the Power 520 will essentially span a wider range at the low end of the product line. It will be supplemented by a new single-socket blade, the JS12. With the increased acceptance of blades and the availability of a new SMB-oriented BladeCenter chassis, the BladeCenter S, more and more customers view blades as a viable alternative to rackmount servers—even for relatively small installations. Product lines tend to naturally get more cluttered over time in response to competitive announcements and customer requests, but IBM has made an honest effort to de-clutter the Power lineup.

From a macro perspective, however, the biggest shift is that there are no longer separate “System p” and “System i” servers. IBM has been steadily tearing down the wall between the two lines. Both i and p have used the same Power Architecture processor for several years now. However, this announcement represents the first time that the wall is truly gone. System i is System p, and vice versa. They’re now both just Power Systems, sharing the same SKUs, the same price lists, the same options.

More cosmetic cleanup extends to nomenclature that had likewise accumulated cruft. Much of it probably “seemed like a good idea at the time” (and perhaps was). Take AIX 5L, a name introduced in 2001 that incorporated a new major version number (“5”) with L (which stood for Linux affinity).[1] Linux affinity referred to a number of code compatibility upgrades that IBM made to AIX. But it was at least as much a strategy statement about IBM’s embrace of Linux as a major company initiative as it was about any low-level technical details. Thus, we had a “5” that tends to be associated with the POWER5 processor when used elsewhere in Power systems nomenclature, and an “L” that seemed a bit anachronistic—a throwback to a time when Linux was new and cool rather than just part of the landscape.[2] So IBM has gone ahead and lopped them off the AIX name.

i5/OS got a similar makeover. It changed from “OS/400” to “i5/OS” when the POWER5 processor started to roll out in the iSeries, then System i. (So, yes, in this case 5 did refer to POWER5!) Now it goes by plain old IBM i. If that seems maybe a bit too short and sweet, remember that Cisco’s network operating system is called IOS and you might see the problem. (Apparently IBM even polled its user base about going with “OS/400” in a sort of retro play and was told in no uncertain terms “NO!”).

PowerVM is IBM’s latest effort to pull together and talk about its virtualization offerings in a systematic way. “Virtualization Engine” was one earlier, less than wholly successful, effort.[3] If anything, it’s become more important for IBM to better organize its virtualization portfolio. Even as virtualization becomes more widespread and important to more businesses, it’s also become more complicated. IBM, for example, has added the ability to move a running LPAR (i.e. Logical PARtition—effectively a virtual machine) from one server to another without shutting it down. Live Partition Mobility is IBM’s answer to VMware’s VMotion. IBM has also supplemented LPARs with AIX Workload Partitions (WPARs, i.e., operating system containers) that can also be dynamically moved, though in more a shutdown, then quickly restart way.

Finally, IBM is taking a stab at rationalizing its Power Systems software as a whole, thus the new Power Systems Software group. In addition to virtualization and operating system/integration layers, IBM groups the software into availability, security, energy (mostly under the EnergyScale umbrella), and management (primarily IBM Systems Director) buckets. Some early changes are mostly about creating a more consistent nomenclature. For example, IBM’s long-standing Unix clustering product, High Availability Cluster Multi-Processing (HACMP), becomes PowerHA for AIX and PowerHA for Linux while System i High Availability Clusters (HASM) become PowerHA for i. However, we expect continued efforts to better align development roadmaps and share technologies as well.[4]

---

1 AIX 5L also has a version number like other OSs do, so the first release of the rebadged AIX was officially AIX 5L 5.1.

2 HP similarly bore HP-UX 11i in 2000 to send the message that it was hip to that cool Internet thing.

3 See our POWER By The Piece.

4 One interesting comment we've heard several times is essentially "Hey, it turns out that i guys have some cool stuff. Now that we're together, we'd like to learn from that." This is a sign that the teams are deeply integrating, not just doing so on the surface as a marketing exercise.

The Rackmount Lineup

The mainstream commercial Power Systems family running on the POWER6 processor now consists of the Power 520, Power 550, Power 570, and Power 595—plus two blade servers, the BladeCenter JS12 and JS22. (See Table 1.)

Also under the Power Systems umbrella are several "specialty" servers such as the Power 575 that is positioned as a heavyweight compute node, whether employed for pure technical high performance computing (HPC) or for more business-oriented analytic tasks. Blue Gene tilts further toward the pure science direction with a huge number of small PowerPC-based nodes running Linux. The QS21 Cell blade, built around the Cell Broadband Engine initially developed for game consoles, can be used as a component of a "hybrid cluster" in which Cell handles the heavy-duty floating point while a more general purpose processor handles other tasks. Because these are specialized HPC products, we won't discuss them further in this note.5

The Power 520 Express,6 to give the system its full name, is the new entry-point to IBM's rackmount Power lineup. The 520 comes with between one and four POWER6 cores and is pitched for a distributed application server, a small database or consolidation server, or—in the case of the i Edition—a complete business solution with an integrated database and application server. The Power 550, which can be configured with between two and eight POWER6 cores, targets a similar set of uses—albeit at higher scale points. Both servers are available in AIX, Linux, and i Editions (i.e. base operating system configurations) but can run any combination of OSs using virtualization.

The modular midrange System p 570 was the first IBM server to get decked out with the POWER6.7 The Power 570 unifies this server and the System i 570. There's no charge to upgrade to the new flavor.

The Power 570 is constructed from up to four discrete building block "drawers," each with two dual-core POWER6 processors. Multiple drawers can be interconnected by cables to create a larger unified SMP server. The advantage of this modular approach is that it lets customers "pay as they grow," incrementally purchasing capacity as they need it, rather than up-front. It's not a new idea. The high end of IBM's System x lineup takes a similar approach in the x86 realm.8 Antecedents go back to Data General and Sequent in the mid-1990s.

Although it has a surface similarity to these other designs, the 570 is cut from different cloth. With the 570, calls to remote memory take only about 25 to 50 percent longer than the best-case local memory access; the exact number depends on other factors, including the number of building blocks installed. This isn't quite a "flat" or purely uniform memory access time—but it's as good as, or better than, many systems using hardwired, inside-the-box interconnects, and significantly superior in both relative and absolute terms to any other modular design on the market.9

The Power 595 is the true “Big Iron” of the lineup, the one that gets to trumpet world-beating benchmark prowess. It can be configured with up to 64 cores running at up to 5 GHz (a new top-end speed grade for POWER6 with this announcement) via 8-core processor books. Each book has 16 to 32 DIMMs per processor book; that’s up to a whopping 4 TB of 533 MHz or 1 TB of 667 MHz DDR2 in a fully-configured Power 595 system.

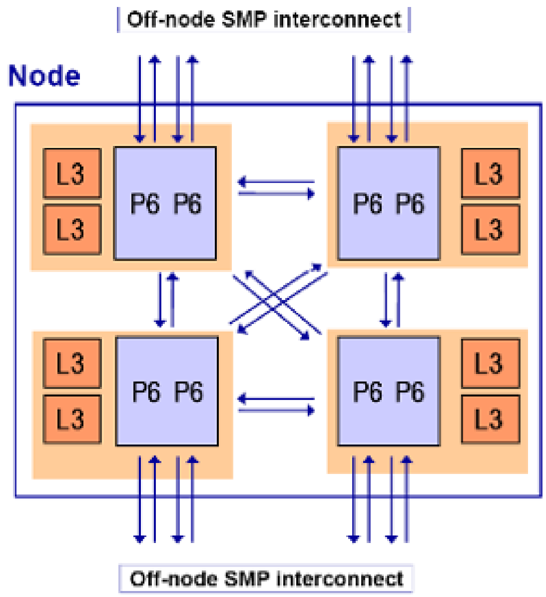

The Power 595 is constructed out of building blocks consisting of four POWER6 chips with their associated L3 cache fully connected into an 8-core node. (In the past, IBM has packaged such nodes into a multi-chip nodule (MCM) but not in this case.) Up to 8 nodes are then fully connected to form up to a 64-core SMP system. A nice feature of this connection topology is that total bandwidth increases faster than compute capacity. Of course, coordination traffic increases as well. However, IBM says that its tests show that even a 64-core Power 595 will saturate memory bandwidth before it will bottleneck on the communications associated with keeping memory accesses coherent (as can sometimes happen on larger SMP servers).

This quantity of CPU and memory hardware (and equally hefty I/O capabilities) translates into equally impressive performance metrics at the system level. IBM has often capped its high-end systems announcements with new TPC-C high-water marks that purport to model some improbably large OLTP environment.10 This time it’s turned in a two-tier SAP Sales and Distribution (SD) Standard Application Benchmark that bests the next closest number of simulated benchmark users—a December 2006 HP Integrity Superdome result—by about 11 percent.11

However, these days the point of such large servers isn’t so much merely to handle mammoth workloads. That a system like the Power 595 is more focused on doing many things well (rather than just one thing) reflects that system horsepower has grown faster than the ability, need, and desire of single applications to use it. Sure, some big databases, ERP, and technical apps still need the biggest boxes available when a task requires too much coordination to run efficiently on a more distributed cluster of smaller systems. But, more and more, big systems are primarily about consolidating tasks in one place, where they can be controlled from a single point and their resources resized as the business needs of the moment dictate.

---

5 See our *Blue Gene*s Teraflop Attack, Cell Trickles Up to Supercomputing, and Power 575: Radical and Extreme for detailed treatments of these systems.

6 IBM has used the “Express” term to connote integrated or pre-packaged offerings.

7 See our It Slices, It Dices, It Runs a Lot Faster for a more detailed rundown. The name's changed, but the product is the same.

8 See our *IBM*s X3 Heads Into the Yukon.

9 NUMA (Non-Uniform Memory Access) describes any system in which communications from one processor to memory take longer than performing the same action on some other processor or some other part of memory. IBM objects to the "NUMA" term because historical examples could take more than ten times longer to access memory located on a remote building block than to access local memory; thus, NUMA implies "slow" in some peoples' minds.

10 See our Another Day, Another Million tpmC and TPC-C Passes Escape Velocity.

11 IBM hits its numbers with half the number of cores. With ISVs such as Oracle that charge per core, this can lead to lower software licensing costs. In other cases such as SAP that are user-based, it makes no big difference. Unlike TPC benchmarks, SAP SD does not include a pricing component, although IBM claims the price of the two systems is "comparable."

The Blades

With the POWER6 generation, IBM has better aligned its Power Architecture blades with its rack- and frame-oriented lineup. Whereas the previous JS21 blade was built around the PowerPC processor, the quad-core BladeCenter JS22 Express and the new dual-core BladeCenter JS12 Express use the same POWER6 processor as the rest of the servers.

The JS21 had some significant wins in the high performance computing arena (such as MareNostrum at the Barcelona Supercomputing Center), but it only saw significant commercial use in fairly narrow niches. For example, some of the retail applications that IBM deploys run on AIX (while others run on Linux or Windows), so having a Power blade lets IBM integrate applications running on disparate architectures and operating systems in a single chassis, an approach the company calls "Store in a Box."

The JS22 has seen broader commercial use. In part this reflects that the JS22 just has more processor oomph. This makes it better suited for workload consolidation using virtualization—a general Power Systems strength and target—as well as capable of handling heavier-weight business applications. It also simply reflects that there is greater commonality with the rackmount server line that is often used in conjunction with blade servers. The reality is that IT shops are far less tolerant of deploying what they consider "one-offs" that are in some way not part of a mainstream offering. The addition of POWER6 helps to more clearly cement Power blades as a standard part of the product line.

With the latest announcement, IBM has also moved down-market with its new single-socket JS12 offering. This reflects the major midmarket and distributed site push that IBM is making with BladeCenter. This is reflected on a couple fronts in addition to the blade itself.

One is the addition of the i operating system to blades. To appreciate this in context, it's worth remembering that the value proposition of i has often been the integration of disparate applications at midmarket companies. This included not only various business apps written for i5/OS but even Windows applications that ran on an Integrated xSeries Adapter (IXA)—basically, an x86 server lashed to an iSeries server. (i5/OS V5R4 also added the ability for x86-based IBM servers to integrate with System i storage using an iSCSI connection.)

IXA was never an ideal solution however. It was a special piece of hardware and tended to lag the technology in standard off-the-shelf x86 servers. Blades provide an alternative, and now more familiar, path to such heterogeneous integration. One can run i, AIX, and Linux applications on JS21 or JS22 blades—using PowerVM, even on the same blade, in different LPARs. Windows and x86 Linux apps can run, in the same BladeCenter chassis, on standard Intel- or AMD-powered blades.

BladeCenter really has become the integration point, not just for various x88-oriented workloads, but now equally for environments that want or need to run on POWER. The BladeCenter S chassis is optimized for midmarket and smaller customers, who tend to use smaller configurations, and often locate their IT outside datacenters. Integration through blade infrastructure should be a great match for the same midmarket customers for whom System i and the AS/400 before it, has held such appeal.

POWER6

POWER6 boosts performance considerably over both the PowerPC and the POWER5+.12 Much higher processor frequency is part of the story; the Power 595 can be configured with up to a 5 GHz speed grade—a new high-water mark for all mainstream CMOS processors. It also has particularly sophisticated Symmetrical Multi-Threading (SMT) that makes each physical core appear to the operating system as two logical CPUs. The goal is to keep execution units in a processor as busy (and therefore as productive) as possible by giving them more than one thread to churn on. POWER6 tracks how threads use shared resources like cache slots and Global Completion Table entries, and adjusts their allocations accordingly.

IBM has also added specialized execution units to accelerate specific types of operations, and further augmented the chip’s reliability features. Notable changes in this area include:

AltiVec (VMX) vector instructions. POWER6 inherits from the PowerPC the AltiVec (VMX) vector instructions that it uses to accelerate floating-point code. This unit is a big part of the Power’s success in HPC applications.

Decimal Floating Point. Another application-specific execution unit accelerates Decimal Floating Point operations. This one’s for the commercial folks who have to work with dollars and cents using a format called Binary Coded Decimal (BCD). This format makes calculations exactly—an important thing when dealing with money. The problem has been that the associated software algorithms are relatively expensive of computer time at the huge scales they’re performed at businesses and financial institutions. The POWER6 accelerates these operations in hardware.

Instruction Retry. Beyond performance, POWER6 also brings in a venerable mainframe reliability feature-instruction retry, dubbed “processor recovery” in this incarnation. The retry operation can occur on a different processor from the one on which the instruction was originally executed. Thus, in the event of a hard error, another processor (such as a Capacity on Demand reserve processor) can substitute for the failed one.

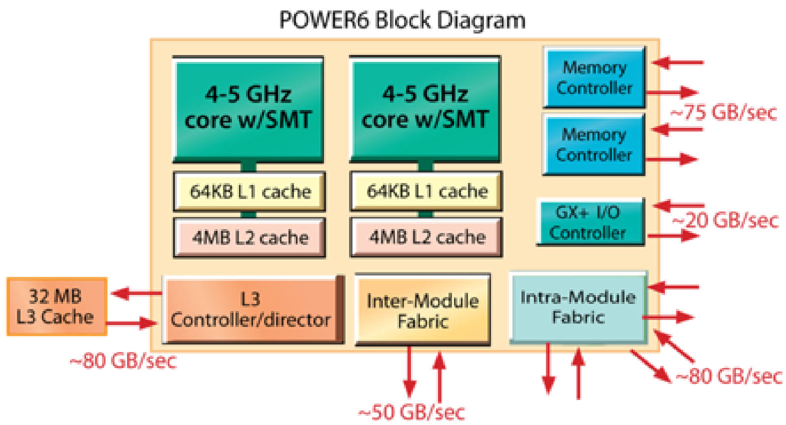

Especially from the perspective of the larger systems in the family, however, just as significant as the features within the core is the way that the cores communicate with each other and with the rest of the system. As with POWER5, POWER6 processors have integrated memory controllers—a much-touted feature in AMD's Opteron and an approach that Intel is moving to with its QuickPath Interconnect. This combined approach has consistently delivered very fast memory access times—a key application performance ingredient—even with memory that is physically distant from the processor making the request.

One significant POWER6 architectural change is a new two-tier interconnect architecture and coherency protocol in addition to other changes aimed at increasing bandwidth and lowering latency. For example, POWER6's mechanisms that make sure that the processor caches are coherent with main memory added an advanced heuristic that has knowledge of the way that memory is allocated and the way that memory is configured. This allows some actions to be kept local to a group of processors and memory and thereby avoid clogging remote parts of the system with unnecessary traffic.13 (Thus, in a Power 595, coherence traffic can often be kept local to an 8-core node so it doesn't need to go onto the off-node interconnect.)

---

12 See our POWER6 Cranks the Clock for a more detailed description of POWER6 features and underlying technology.

13 There are high-level conceptual similarities to the snoop filter in IBM's X4 chipset although the details are quite different. See our *IBMs* Uniquely Scalable Xeons.

PowerVM

One major virtualization trend that we've observed over the past couple of years is that even companies with very deep capabilities in some particular approach have come to the realization that one size truly does not fit all. Thus Sun went from trumpeting hardware partitions (and then later Solaris 10 containers) as the ultimate answer to Life, the Universe, and Everything to something that's part of a toolkit. From the client side, we've seen Citrix evolve from all-Presentation-Server-all-the-time to a much broader application delivery strategy. And we've seen IBM go from "LPARs is the answer. What's your question?" to matching a variety of technologies to differing sets of requirements.14

One example of this is the bright spotlight that IBM now shines on z/VM, long a red-haired stepchild compared to its z/OS sibling in the mainframe space.15 However, it's also reflected in the complementary set of virtualization capabilities that IBM now offers on its Power line.

The first is PowerVM—roughly speaking, the set of functions that used to go by "Advanced Power Virtualization," plus some recent enhancements. (As in the case of PowerHA, PowerVM is both an umbrella term for a related set of capabilities, and part of the name for specific products.)

The foundations of PowerVM are micro-partitions, as many as ten per POWER6 core up to a maximum of 254 per server.16 Relative to the largely software-based approaches of the x86 world,17 the controlling hypervisor is in firmware and tightly ingrained at the lowest levels of the system. The virtualization control logic works with processors, I/O cards, service processors, operating systems, systems management tools, and other components to coordinate the allocation of system resources.

Reflecting its Big Iron roots, IBM's approach to partitioning and virtualization initially focused on optimizing and consolidating workloads within a single server. However, with the increased buzz around what is often called "dynamic infrastructure" or similar names, IBM has shifted its attention to virtualization that spans multiple machines, and to integrating virtualization with management tools such as its Systems Director.

With POWER6, IBM introduced Live Partition Mobility, the ability to move a running LPAR from one physical server to another. IBM says the observed "pause" for the shift is about two seconds—within the TCP/IP timeout window, so that network clients won't even notice that the LPAR is suddenly running on another server. Although software-based approaches such as VMware make similar claims, their reality is highly load dependent. (We've heard anecdotal x86 user stories suggesting that, when there are a lot of memory writes going on, it can be difficult to complete the transfer.) To use Partition Mobility, the servers must all be on the same network subnet, and all of the I/O resources in the LPAR must be virtualized through a VIOS (Virtual I/O Server). Unlike "pure" VM approaches, IBM allows I/O resources to be either virtualized or, for physical devices, to be exclusively controlled by an LPAR; but for purposes of mobility, only the first option works.

PowerVM comes in Express, Standard, and Enterprise Editions. Express is the free "try it, you'll like it!" offering; it's limited to 3 total LPARs. Standard enables the full 10 LPARs/core level of virtualization, and adds more advanced functions like Multiple Shared Processor Pools—the balancing of processing power among partitions assigned to shared pools. Enterprise adds Live Partition Mobility as well.

As of AIX 6.1, IBM also has Live Application Mobility,18 the ability to move workload partitions (WPAR) from an instance of AIX running on one server to an instance running on another. The transfer is performed using the Workload Partitions Manager (WPM), a new tool that integrates IBM's previous WLM (workload manager) with the fruits of its Meiosys acquisition. WPM provides a single graphical console for managing system and application WPARs across systems, including creating and removing them, starting and stopping them, and relocating them. Mechanically, the transfer works by checkpointing the application to disk and then restarting it on a different server; the application is frozen while the transfer takes place (with the result that network connections to that application would likely timeout and drop while the transfer is taking place). Although Live Application Mobility isn't as flexible as Live Partition Mobility, neither HP nor Sun offer the ability to easilymove workload containers.

The WPAR style of virtualization—often called containers or operating system virtualization—has proven popular for uses where minimizing overhead is a priority. Hosting providers are the canonical example of this use case, although we've also spoken with enterprise customers who have likewise gone the container route to minimize their use of hardware resources for virtualization.19 For example, with

WPARs there's only one copy of the operating system for the entire server or LPAR, no matter how many containers you fire up. For the same reason, you only have to patch that one copy of the operating system. The tradeoff is that you can't mix operating system types or versions using WPARs, don't get the same degree of fault isolation, and don't get the same heavy-duty, under-high-load ability to transparently transfer a running workload from one server to another.

Finally, IBM offers PowerVM Lx86. Formerly called the IBM System p Application Virtualization Environment (System p AVE or, informally, pAVE), it uses Transitive's QuickTransit to let 32-bit x86 Linux applications run atop Linux on a Power system.[20] IBM's interest in running Linux on its Power-based servers goes back a few years, but, for a variety of reasons, it never really hit its stride. Now, though, IBM is leveraging LPARs to promote the platform for Linux server consolidation. If consolidating a web farm, for example, most of the applications—such as the Apache Web server—would likely have corresponding native Power Linux binaries available. However, there will often be some Lx86 code or third-party applications that don't have Power equivalents. That's where Lx86 comes in—providing a way to run these "completers" as dynamically translated Linux-on-x86 applications.

---

14 HP, by contrast, long has gone for the smorgasbord approach, unified under the VSE (Virtual Server Environment) umbrella. See our *HP*s VSE Shows the Practical Side of Server Virtualization.

15 See our z/VM: Teddy Bears and Penguins.

16 Micro-partitions sometimes also go by the term Dynamic LPARs (DLPARs) because they are descended from the logical partitions that IBM repurposed from the mainframe for its Power line.

17 Chip technologies such as Intel-VT and AMD-V provide some assistance but all the core hypervisor functions reside in higher-level software. POWER6 also includes instructions and modes to support virtualization.

18 Not to be confused with Live Partition Mobility, which applies to LPARs and works with multiple operating systems (AIX and Linux).

19 See our The Server Virtualization Bazaar, Circa 2007 for more background on the different styles of virtualization.

20 See our *Transitive*s Translations.

Conclusion

As IBM never fails to remind those who want to listen—and, indeed, those who do not—Power has been making great revenue growth and market share strides over the past few years. Indeed, Power, especially running AIX, has been the most consistent star of the Systems and Technology Group firmament. That's no small feat given that IBM has quite a stable of strong and successful products. It's quite the turnaround. Set the Wayback Machine to the mid-1990s and it wasn't clear how serious IBM was about even staying in the RISC processor game.That was then. This is now.The Power landscape can't be a particularly pleasant one for competitors to view. That's not to say that the likes of HP and Sun don't have their own strongholds and their own competitive counters. But the strength of the Power lineup overall, both its servers and its supporting software, certainly doesn't lend itself well to direct frontal assault. Thus, for example, we see Sun's focus on running Solaris on x86 and focusing its UltraSPARC efforts on far more aggressively multithreaded designs.[21]

The task of competitors isn't made any easier by the fact that IBM has cribbed and incorporated so many of their plays. IBM's unified Power, combined with virtualization, is cut from a similar cloth as the Multi-OS on Itanium strategy that HP once promoted so strongly. And rather than continuing to disparage the operating system containers that were the arrowhead of Sun's virtualization strategy, IBM developed its own flavor—and then one-upped Sun by adding the ability to migrate them from one server to another.

Yes, the merging of i and p and the related organizational changes are no easy tasks. As a product manager in a past life, I look at the vast roto-tilling of names, trademarks, product and option numbers, order channels, and backroom procedures evident in these recent announcements, and I can only imagine the screams of "But it can't be done that way!" that must have echoed within the walls of IBM's Austin, Somers, Poughkeepsie, and Rochester locations.

But the results in evidence are salutary. Power Systems now sits atop the Unix hill, planning how to add to the lands under its domain.

---

21 See our *Sun*s Vic Falls: Two is Better than One.

Table 1. The Power Systems lineup

	BladeCenter JS12 Express	BladeCenter JS22 Express	Power 520 Express	Power 550 Express	Power 570	Power 575	Power 595
System package	BladeCenter	BladeCenter	4U, 19" rack/tower	4U, 19" rack/tower	4U (per block), 19" rack	2U, 24" system frame	4U, 24" system frame
POWER6 cores	2	4	1,2,4	2,4,6,8	2,4,8,12,16	32	8 to 64
Frequency (GHz)	3.8	4.0	4.2	3.5, 4.2	3.5, 4.2, 4.7	4.7	4.2, 5.0
Max memory (GB)	64	32	64	256	768	256	4096
Max PCI Slots	1 PCie 8x 1 PCI-X	1 PCie 8x 1 PCI-X	3 PCie 8x 2 PCI-X	3 PCie 8x 2 PCI-X	16 PCie 8x 8 PCI-X	4 PCie 8x	20 PCI-X
Max PCI Slots w/ I/O drawers	0	0	2 PCie 8x 58 PCI-X	1 PCie 8x 58 PCI-X	12 PCie 8x 200 PCI-X	4 PCie 8x 20 PCI-X	600 PCI-X
Redundant hot-plug power	Optional at chassis level		Optional		Standard	Standard in system frame	
PowerVM Express Edition	No	No	Optional	Optional	No	No	No
PowerVM Standard Edition	Standard	Standard	Optional	Optional	Optional	Optional	Optional
PowerVM Enterprise Edition	Optional	Optional	Optional	Optional	Optional	Optional	Optional
Operating system support	AIX, i, Linux	AIX, i, Linux	AIX, i, Linux	AIX, i, Linux	AIX, i, Linux	AIX, Linux	AIX, i, Linux