April 25, 2008

Why Consolidate Linux On IBM System P?

by Brad Day
with Simon Yates and Rachel Batiancila

EXECUTIVE SUMMARY

With the evolution of IBM's Power Systems server (formerly known as System p), the core differentiation of its feature/functionality has first and foremost centered on its sustained advantage in scaleable availability and scaleable performance versus both distributed x86 and other RISC and Itanium 2 competitors. However, System p can really no longer be called a Unix/RISC enterprise computing platform. IBM has now converged the System p and System i technologies into a single Power Systems Group, and its re-branded PowerVM technology is taking center stage as the "scaleable virtualization platform for your mission-critical Unix, Linux, and i5/OS applications." Using System p for running Linux workloads pulls on three core drivers: 1) Linux's ability to take advantage of the POWER6-based overall performance scalability; 2) its extreme reliability and availability features; and, most importantly, 3) System p's ability to benefit from PowerVM's affinity for consolidating and optimizing Linux workloads.

QUESTIONS

1. Why should I deploy Linux on IBM's System p?
2. What core benefits does PowerVM bring to Linux on POWER?
3. What unique LPAR advantages in PowerVM can benefit Linux on POWER?
4. What are micro-partitions and shared-processor LPARS; what do they bring to the Linux environment?
5. How do Linux workloads use PowerVM's dynamic reconfiguration, virtual LAN, and virtual I/O functionalities?
6. Can all the same RAS functionalities available to AIX users equally support Linux?
7. What is PowerVM's Live Partition Mobility, and can Linux take advantage of it?
8. What Is PowerVM Lx86, and how does it keep the Linux ISV ecosystem healthy?

POWERVM'S TRACTION WILL DEFINE THE FUTURE SUCCESS OF LINUX ON POWER SYSTEMS

In November 2007, IBM formally released AIX 6, which formed the foundation of its decision to create a Power Systems software organization. While AIX 6 created new and enhanced functionality in the areas of management, energy efficiency, security, and availability, PowerVM virtualization breakthroughs — such as live applications and partitions mobility — garnered most of the System p

customer attention. IBM also completed its systems transition to POWER6 — bringing 1.5x to 2.5x performance improvements to server models ranging from BladeCenter J22 to its entry server (System p 520) and up through midrange server models (System p 550). Driven by PowerVM's virtualization improvements — coupled with these System p blade, entry, and midrange competitive x86 alternative solutions — customers are starting to investigate both the technology and economic advantages of transitioning and consolidating Linux/86 workloads on a System p alternative.

Before considering this transition, however, due diligence requires an understanding of the value that the broad and deep PowerVM virtualization stack can bring in support of Linux and its associated applications workloads.

1. Why should I deploy Linux on IBM's System p?

If you asked System p customers to tell you the top three reasons why they chose System p in the first place, they would probably tell you that it was performance, scalability, and reliability. Linux on System p gives you those same enterprise computing advantages while allowing you to run the Linux operating system and applications like Web services, line-of-business applications, and infrastructure services. Linux for System p is designed for Linux solutions that require a full 64-bit architecture or the high-performance capabilities of the POWER6 processor.

Approximately 3,000 Linux applications have been ported through IBM's Chiphopper program to support System p natively. IBM's PowerVM Lx86 technology (formerly called Pave) is bundled into every PowerVM edition (Express, Standard, and Enterprise).[1] It is this unique software translation technology that will help provide the ongoing healthy ramp of Linux workload adoption on System p.

2. What core benefits does PowerVM bring to Linux on POWER?

PowerVM virtualization allows you to run one or more instances of Linux along with your existing AIX 5L operating systems. With this capability, you can consolidate workloads from several separate servers onto a single System p, potentially increasing overall systems utilization levels. Consolidating existing distributed x86 Linux environments into a System p virtualization environment can also help address the challenges of reducing space, power, and overall life-cycle costs.[2]

The PowerVM stack improves business responsiveness and operational speed by dynamically re-allocating resources to applications as needed to better match changing business cycles or handle unexpected surges in user demand. It can reduce life-cycle costs and make better use of IT assets by significantly improving server use and the sharing of I/O resources. Lastly, it can simplify infrastructure management by making workloads independent of the hardware resources they run on, enabling you to make business-driven policies to deliver resources based on time, cost, and service-level requirements.

3. What unique LPAR advantages in PowerVM can benefit Linux on POWER?

PowerVM can provide low-cost application security and isolation with EAL4+ certification, supporting partition stability, security, and fault isolation due to the POWER hypervisor implementation. In addition, PowerVM can allow for far better granularity than hardware partitioning techniques, as your partitions don't have to be carved in increments of systems boards. LPARs can provide complete security isolation in both the kernel and applications spaces; in other solutions, the kernel level is still exposed when you run in an applications-container-type solution. Another powerful feature is the fact that there isn't a limit on how large your logical partitions can be. LPARs can grow to the full size of the physical machine, which could be important for test LPARs — they could mimic a real production environment. In addition, System p memory can be dynamically reallocated; System p I/O can be virtualized and is reconfigurable; and a mix of dedicated and shared I/O is supported across partitions.

4. What are micro-partitions and shared-processor LPARS? What do they bring to the Linux environment?

Shared-processor logical partitions (LPARs) allow you to slice up a System p machine into virtual Linux partitions and give you the flexibility to dynamically change the allocation of systems resources for those environments. The micro-partitioning feature on System p servers provides the capability to create multiple virtual Linux partitions within a processor to a granularity of one-100th of a CPU, with one-tenth of a CPU minimum partition. In addition, any of the virtual Linux servers may run on any of the physical processors. This means that the processor of Linux resources is fully shared, making it possible to run the physical server at very high utilization rates.

5. How do Linux workloads use PowerVM's dynamic reconfiguration, virtual LAN, and virtual I/O functionalities?

It is possible to move Linux system resources, physical processors, virtual processors, memory, and I/O slots dynamically between partitions without rebooting. As a function of the POWER hypervisor, the virtual LAN allows secure communication between Linux — or Linux and AIX — logical partitions without the need for a physical I/O adapter. The ability to share Ethernet bandwidth securely across multiple partitions increases hardware utilization substantially. Virtual I/O provides the capability for a single physical I/O adapter and disk to be used by multiple Linux and/or Linux/AIX logical partitions for the same server, which essentially allows you to consolidate I/O resources and minimizes the number of I/O adapters you require.

6. Can all the same RAS functionalities available to AIX users equally support Linux?

As individual System p systems become capable of hosting more software images for server consolidation, the importance of predicting, isolating, and handling outages becomes more significant. Hardware and operating system functions have been integrated into the System p design to: monitor systems operations; predict where outages might occur; isolate outages that do occur; handle the outage condition; and, when possible, continue operations. System p servers now support

high levels of concurrent error detection, fault isolation, recovery, and availability. System p has comprehensive RAS feature/functionality and can equally support a Linux applications workload environment. It includes: Chipkill and ECC memory; disk mirroring; journaled file systems; redundant, hot-plug power and cooling; error log analysis; boot-time processor and memory de-allocation; first-failure data capture; hot-swapping of disk drive; dynamic processor de-allocation; and hot-plug PCI disks.

7. What is PowerVM's Live Partition Mobility, and can Linux take advantage of it?

One of the unique features in PowerVM Enterprise Edition is Live Partition Mobility. It allows clients to move a running partition from one physical POWER6-based System p to another POWER6-based System p without any application downtime. This helps avoid any application disruption for planned system maintenance, provisioning, and/or workload management. In support of any running Linux workload, the migration operation just takes a few seconds and maintains complete transactional integrity. The migration actually transfers the entire systems environment, including processor state, memory, attached virtual devices, and connected users.

You can also use Live Partition Mobility for server consolidation. If you have partitions with workloads that have widely fluctuating resource requirements over time — such as an end-of-the-week payroll run — you could use Live Partition Mobility to consolidate partitions to a single server during an off-peak period, allowing you to power off unused servers just prior to the peak period. This approach can offer energy savings by reducing the number of System p POWER6 cores running at any given time.

8. What Is PowerVM Lx86, and how does it keep the Linux ISV ecosystem healthy?

IBM still encourages ISVs to perform the code porting of their applications from Linux/x86 to the native version of Linux on POWER through the Chiphopper set of porting tools. However, IBM's PowerVM Lx86 provides a just-in-time and productive alternative solution: You (or an ISV) can take the original installation media of a Linux on x86 application and install it "as is" with a Linux on POWER partition running on System p. PowerVM x86 provides a compelling environment in which the benefits of running on System p can be made available to customers who need to run applications that are currently only available on x86 platforms. Most importantly, you don't need any application changes to get the same high degree of reliability, availability, and serviceability (RAS) — or the same PowerVM functionality that native Linux and/or AIX environments use.

The use of PowerVM x86 to run x86 applications on POWER systems is transparent to the applications. The objective for x86 applications is that they just run. You can install most 32-bit x86 binaries in the same way that you would install them on an x86 system. As this is a translation technology, these applications are just executed via a command that starts an application. The x86 application binary is ready to run without the technical effort or the delay of a native port. As an aside, Forrester clients report that the more consumptive line-of-business Linux/86 applications —

when transitioned and running on POWER6 System p — have offered performance scalability results ranging from approximately 70% to 85% of what a similar applications workload outcome would be as a System p native port. In short, this translation technology is very efficient and the results are still compelling — as 70% to 85% System p workload performance results (on a core versus core basis) still fully outperform an equivalent Linux/x86 result.

ENDNOTES

[1] After 18 months of collaboration between IBM Labs, software engineering, and product teams and business partner management, IBM formally released the IBM eServer Application Advantage for Linux — aptly code-named the "Chiphopper" program. Through the Chiphopper offering, ISVs that typically developed their application for Linux/x86 server customers will be able to support a single version of their existing Linux application across IBM's entire eServer platform offerings. The Chiphopper program will increase their customer market reach and revenue potential for their applications with only a minimal increase in ISV expense. The most compelling aspect of Chiphopper is the significant cost reduction gained by any Linux ISV that wants to migrate its existing Linux/x86 application to a broader set of IBM eServer systems targets and middleware software platforms. See the April 19, 2005, "IBM Unveils Chiphopper: Linux ISVs Take Notice" report.

[2] Virtualization has wide-ranging effects throughout your IT infrastructure. Even though virtualization has been making technology easier to use, it has become even trendier in the past few years. Why the sudden attention? Recent innovations in server or client virtualization are producing more and more business value than prior innovations in system architecture or programming languages, which were more confined to IT. This report examines the progress being made in virtualized infrastructure and how it is the key to building a more change-ready and agile foundation for your business. See the November 8, 2007, "The Virtualization Imperative" report.