IDC
Analyze the Future

Global Headquarters: 5 Speen Street Framingham, MA 01701 USA P.508.872.8200 F.508.935.4015 www.idc.com

WHITE PAPER

IBM i and Power Systems: Leveraging Virtualization in the Datacenter for Flexibility and Business Continuity

Sponsored by: IBM

Jean S. Bozman
July 2008

EXECUTIVE SUMMARY

Business processes and IT processes are becoming aligned, following years when IT infrastructure was deployed and optimized to meet the needs of separate, and competing, business units across the enterprise. From a business perspective, this approach to leveraging a more unified virtualization platform for the entire enterprise is an important step on the road to a more flexible form of processing in the datacenter. Importantly, it is an approach that supports business agility by allowing new workloads to be deployed more seamlessly, and with less operational cost, as business requirements change, over time.

The IBM i operating environment, running on Power server systems, contributes to the goals of flexibility and business continuity by supporting advanced virtualization capabilities and multiple paths of access to storage resources, which increases availability, and enables ongoing business continuity. Applications running on the IBM i operating environment can easily gain access to sufficient computing resources, as needed — and they can be linked to storage area networks (SANs), in addition to traditional use of virtualized storage within the Power server system. The combination of IBM i and Power server systems is aligned with the major trends in the IT industry: virtualization, consolidation, support for high availability, and support for green IT through efficient use of IT resources.

From an IT perspective, this supports the "pooling" of server and storage resources via virtualization. This approach addresses the needs of enterprise customers that see the demands for applications growing over time, especially during "crunch times" when peak demands force a decision about "rightsizing" an application or when systems have run out of sufficient "headroom" to add resources quickly and conveniently. Given the changes on both the hardware side and the software side of the platform, IBM customers can now add more physical server capacity, even as the software components of the solution gain rapid access to more server and storage resources.

From a business perspective, IBM i provides another approach to virtualizing hardware resources, given its integrated functionality for provisioning and managing workloads — and its ability to tap network resources (such as SANs) to support end-to-end applications running across the entire enterprise. Placing the emphasis on maintaining continuing access to system resources by end users and end customers, IBM's systems address the twin goals of business continuity planning: recovery time objective (RTO), reducing the time to recovery for an outage, whatever the cause

(e.g., network, power, or failure of a hardware or software component), and recovery point objective (RPO), returning the IT system to the state it had prior to the outage as quickly as possible.

INTRODUCTION

The datacenter is being transformed, and businesses are rethinking how best to deploy applications and workloads for the most efficient and reliable operations to reduce operational costs. This transformation process brings with it an opportunity to consolidate workloads onto fewer server footprints, which assists energy-efficiency projects within the business and addresses IT operational costs due to IT staff time for administration, unplanned downtime, and ever-rising power/cooling costs.

IBM Power Systems product line supports a broad range of workloads running on IBM AIX (Unix), IBM i, and Linux operating systems. This paper focuses on the IBM i operating environment and its ability to support business-critical workloads running on Power Systems. IBM i (formerly known as i5/OS) supports virtualized, highly available transaction processing and line-of-business (LOB) applications and is typically deployed in industries such as financial services, manufacturing, and distribution, where business continuity that ensures access to end users and end customers in the organization and the supply chain is especially important.

Business Platforms for Resilient Computing

Customers are looking for computing environments that efficiently deliver data services to promote business processes. In their evaluations of new equipment, they are looking for server platforms and server operating systems that support reliability, availability, security, and manageability. Today, achieving reliability and increasing availability in IT are considered the top priorities for many organizations, and the top challenges in the datacenter, as seen in IDC research studying top concerns of datacenter managers, reflect this agenda for IT systems.

Figure 1 clearly shows that concerns about power/cooling, availability, and the utilization of datacenter "real estate" are top of mind for many IT managers. All of these factors contribute to high operational costs — and must be addressed in next-generation IT deployments.

FIGURE 1

Top Datacenter Challenges

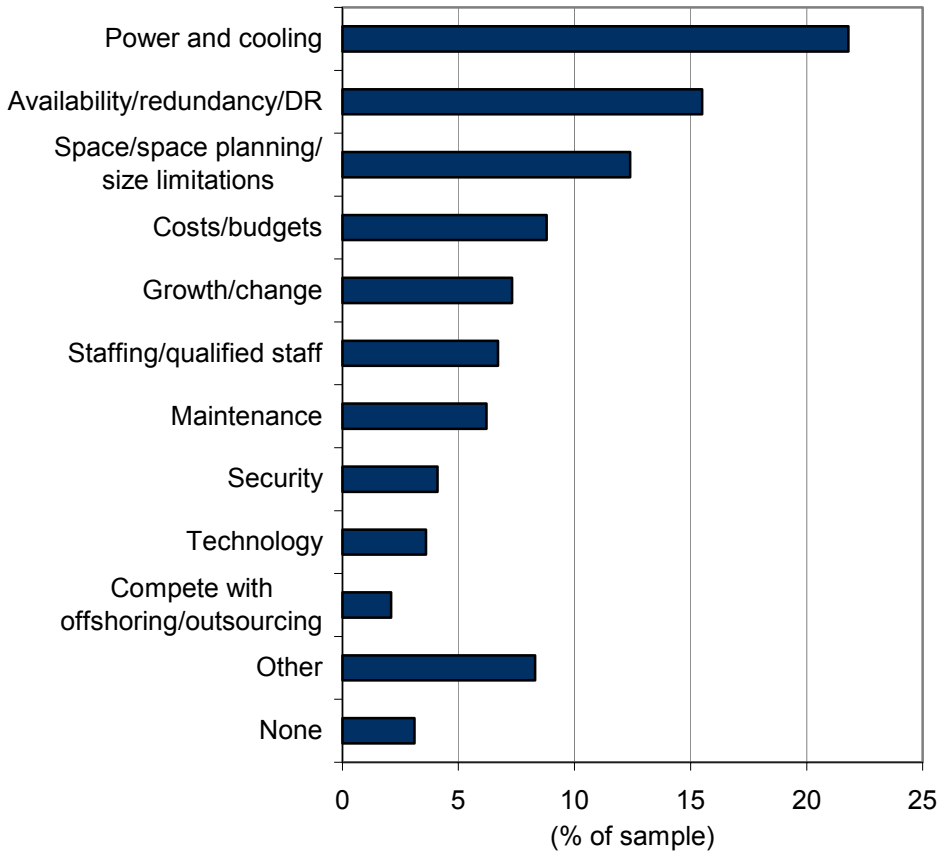

Source: IDC, 2008

To the extent that these top challenges can be addressed in the datacenter, IT can focus more intently on supporting business processes and business applications rather than on the underlying IT hardware itself. After years of optimizing, tuning, and adapting applications to run on specific hardware platforms, there is a desire to better align business needs with IT platforms — and not the other way around.

IT efforts to break down barriers imposed by information "silos" within the datacenter have been hampered by the heterogeneous nature of the servers that reside there — and by the absence of virtualization software that allows workloads to grow nondisruptively on one server or to flow easily from one server to another.

The technical aspects of this platform should be invisible to the business units, which are relying on their ability to access data services on a continuous, or near-continuous, basis. If any service interruptions are caused by power outages or natural or man-made disasters, the IT infrastructure needs to be able to respond flexibly, and to restart processing on alternate resources, in a seamless way. It would be ideal to be able to ignore all of the variations in the hardware when planning a capacity expansion in the datacenter.

Although the process of transforming the datacenter into "pools" of resources that can be tapped as needed will take many years to accomplish, it is possible to make substantial progress by consolidating workloads onto highly virtualized server platforms within the larger datacenter. To the extent that those platforms are highly virtualized — supporting the task of managing many workloads efficiently and high levels of system utilization — the task of making the overall datacenter more efficient will be greatly enhanced. As a result of these considerations, many organizations are moving toward the idea of on-demand systems that automatically add capacity as business needs change — and subtract it when it is no longer needed.

To move to a more dynamic IT environment, customers need to evaluate business processes across the enterprise, to determine which components within the infrastructure need the most IT resources. IDC describes "dynamic IT" as an approach to computing that makes IT more flexible and adaptive while breaking down barriers between what were historically information islands, or silos, within the enterprise — each focused on a separate set of workloads.

Importantly, dynamic IT maps business processes to IT processes, improving the alignment between the two types of processes — and the ability to adjust IT resources to sustain changing business requirements, as needed. This mapping process should take into account the hardware, software, storage, and interconnects that link the IT systems, which assist the business processes. IDC notes that this mapping process must take into account the fact that service-level agreements (SLAs) will vary by business unit and by application. Thus, this is neither a one-size-fits-all situation nor a "rip-and-replace" situation, which would cause unnecessary disruption to the enterprise. Therefore, careful thought and planning should be applied to identifying requirements for each aspect of the IT infrastructure — and to applying appropriate solutions to the systems addressing business processes.

Managing Workloads to Reduce Rising Operational Costs

Efficient computing — and all it brings with it — is the aim of many IT transformation efforts today, largely because the support for business processes has become paramount and the IT infrastructure has become too complex to operate efficiently. Why is this occurring now? Increasing operational costs, often associated with power/cooling, rising energy costs, the need to improve the availability of business-critical applications, and datacenter space requirements, have forced the issue.

Operational costs are rising in today's datacenter, driven by the IT staff costs of managing large numbers of servers and rapidly rising energy costs related to power and cooling of densely packed IT infrastructure. As Figure 2 shows, the worldwide installed base of servers rose dramatically from 1996 to 2006, due to the deployment of large numbers of small servers on a worldwide basis. This trend had the effect of increasing the costs of management, as IT staff focused on provisioning workloads on large collections of small servers — and then updating and upgrading those servers and addressing the increasing complexity of large populations of rack-optimized servers within the limited real estate of traditional datacenters.

FIGURE 2

Worldwide Server Installed Base and Costs for Acquisition, Management, and Power and Cooling

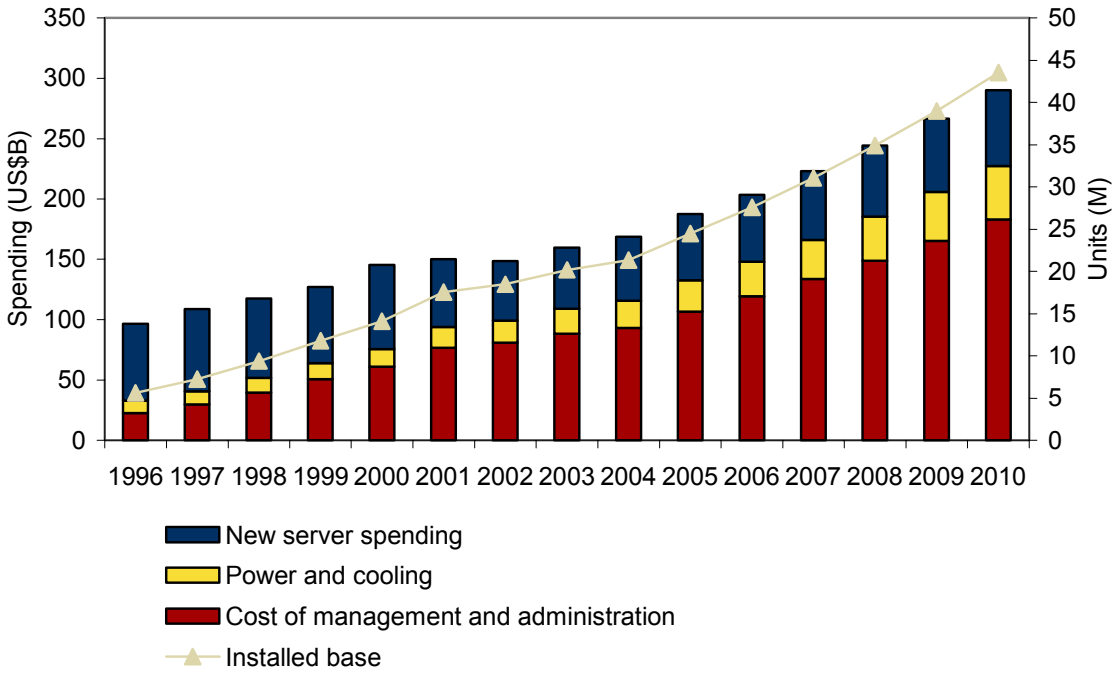

Source: IDC, 2008

In recent years, many datacenters have realized that underutilized servers — many of them small servers in scale-out deployments that were installed to reduce capital costs (capex) — have been impacting overall operational efficiency through increased operational costs (opex). Many datacenters have addressed this rise in opex through a cycle of virtualization of the hardware itself and consolidation of workloads onto a smaller number of server "footprints" in the datacenter. In many cases, though, the number of servers at sites multiplied over time, greatly increasing the installed base of servers that must be operating, managed, and maintained. Figure 2 shows how the worldwide installed base of servers has boomed from roughly 5 million servers in 1996 to more than 30 million servers today.

Today's datacenters are working to reduce operational costs by applying new technologies and new approaches to deploying IT systems. This process of IT infrastructure transformation is still under way, and many sites are seeing improvements in operational efficiency, energy efficiency, and management costs. The aim of this transformation is to reduce IT staff costs while improving the help that business units gain from the applications being run on IT systems.

As enterprises move to group their server resources in new ways to gain greater efficiency in processing and more control of where — and when — applications are processed, new technologies are impacting operational costs, including those associated with power/cooling, IT staff, and overall maintenance. Indeed, efficiency in processing is the goal of many virtualization and consolidation initiatives, which improve server resource utilization by allowing workloads to be sent to the appropriate computing resources, and moved, as demand for processing "peaks" or ebbs, based on business requirements.

IBM Power Systems

IBM has unified its servers based on the POWER processor, making it possible to host multiple operating systems on Power platforms — and to virtualize these systems to support IT flexibility and business agility. In this way, the IBM System i servers and IBM System p servers are now delivered as a single Power Systems product line including blades, volume servers, midrange servers, and high-end servers — all of them based on POWER6 processors.

The new IBM Power Systems product line includes POWER6 processor–based systems and blades — supporting a broad range of capacity from 1-core to 64-core platforms that extends from the Power 520 volume servers (as defined by IDC's price bands, volume servers are servers that are priced less than $25,000) to the Power 595 high-end enterprise models (high-end enterprise servers are servers that are priced at $500,000 or more). By simplifying the hardware offerings — and supporting three operating systems for all models — IBM is increasing customer choice regarding workload deployment across these building-block offerings, which are acquired and deployed based on their price/performance.

As datacenters optimize their IT infrastructure, this unification of IBM Power Systems as one product line offers more choice and more deployment flexibility for the software environments running on the POWER processor. IBM customers now have the ability to assign workloads to operating systems, as needed, based on their functionality and their support for specific applications. Each IBM Power System supports IBM AIX, IBM i, and Linux running in separate logical partitions, or separate "images," without interfering with the other operating systems running on the same machine. The combination of physical partitions (micropartitions) and virtualization (logical partitions) supports the isolation of operating system images. With IBM i, that level of isolation and control is further enhanced through the use of "subsystems," enabling multiple workloads to be managed within a single operating system image. The combination of virtualization techniques — logical partitions and subsystems — allows the workload mix on each server to be optimized for performance, reliability, scalability, and availability. Because the software environment is virtualized, and for that reason is abstracted from the hardware itself, workloads can be moved around the system resources without disrupting the hardware operations or requiring a reboot of the entire system.

PowerVM for Virtualization

IBM provides a hypervisor engine — called PowerVM — that provisions workloads, allowing them to run across POWER processor–based servers and supporting efficient deployments of those workloads and unified management across all of the system's hardware resources. PowerVM is a unifying technology that supports images of IBM operating systems, running side by side, within separate, logically isolated partitions (called LPARs) of an IBM Power System. These operating systems can be mapped into as many as 10 micropartitions within each POWER6 processor.

Figure 3 shows that the PowerVM hypervisor supports three operating environments natively on POWER6 hardware — IBM AIX, IBM i, and Linux — and works with software products from IBM and ISVs to bring new workloads to a growing inventory of applications running on IBM i. In doing so, the PowerVM impacts management of workloads, energy usage, availability levels, integration of operating system functionality, systems software, and also virtualization of all of the IBM Power System workloads, working with orchestration and management software provided by IBM Tivoli.

FIGURE 3

PowerVM Partitioning IBM AIX, IBM i, and Linux on a Single POWER6 System

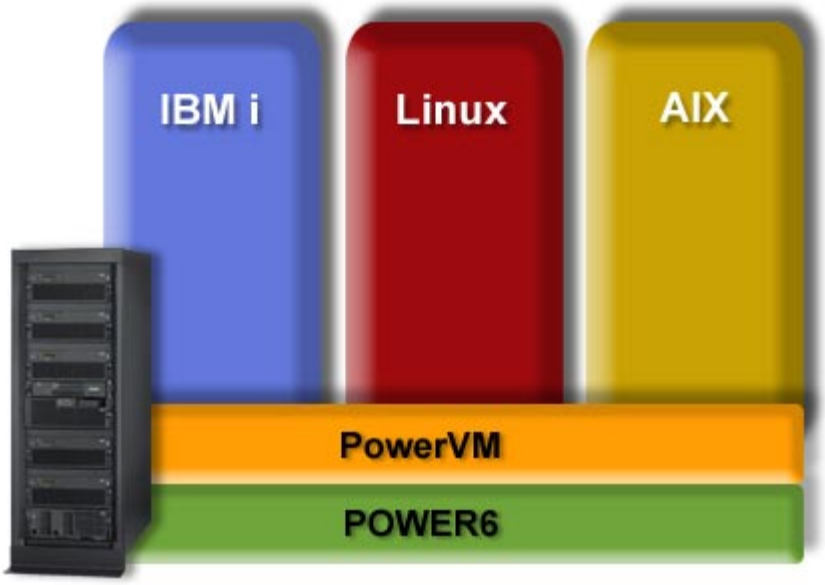

Source: IBM, 2008

In summary, PowerVM allows customers to tap multiple operating systems, and to move appropriate computing resources to workloads, as needed. Because of this, PowerVM supports the goal of dynamic IT in the datacenter, removing barriers that prevented IT from combining workloads that formerly ran on separate servers aimed at hosting one department's applications, rather than supporting workloads across the entire enterprise.

Exploiting SANs: A New Approach to Storage for IBM i in the Enterprise

The latest version of IBM i release 6.1, when combined with the new POWER6 processor–based systems, brings new levels of SAN performance to users of IBM i systems. IBM has stated that its new strategy is to optimize IBM i for SAN performance as an alternative to "in-the-box" or internal system storage, which has been the primary approach taken on the iSeries and System i server platforms for many years. This change helps companies that are implementing a broader enterprisewide strategy for storage by managing all stored data together, regardless of the operating system for the server that transmitted the data to the SAN. This approach is enabled by the ability of all SANs to store data from all types of servers, even though the servers are running a range of different operating systems. Given the move to on-demand computing and dynamic IT — which, as defined by IDC, aligns business processes and IT infrastructure more closely — this approach is another technique to break down the barriers to end-to-end application assistance across all enterprise IT infrastructure.

Moving to SAN storage creates storage "pools" that can be resized, as needed — and this creates more IT flexibility than was possible when storage was contained within a single system. Even though IBM has stated that SAN storage will be the preferred option for future deployments, especially in larger enterprises, longtime System i customers (and customers that installed IBM iSeries and IBM AS/400 servers) can retain their current internal storage without a deadline for changeover.

With the use of micropartitioning on IBM Power Systems, and virtualized storage on SANs, workloads running on IBM i have access to the full capacity of scalable servers, so that rapidly growing workloads can easily gain access to sufficient computing resources, as needed. This addresses the needs of enterprise customers that see demands for applications growing over time and also adds capacity on demand during "crunch times" when peak demands force a decision about "rightsizing" an application on systems that have run out of sufficient "headroom" to add resources quickly and conveniently. Given the changes on both the hardware side and the software side of the platform, IBM customers can now add more physical server capacity, even as the software components of the solution gain rapid access to more server and storage resources.

PowerHA for High Availability

IBM has also expanded the high-availability (HA) offerings with IBM i on the Power Systems platforms, underscoring the need for HA choices for mission-critical and business-critical workloads, especially those running on midrange and high-end servers, such as IBM's Power 570 and Power 595 servers. These scalable servers have the ability to consolidate many workloads, meaning that any downtime would impact large numbers of end users, if it were to occur. That is why HA software, along with reliability, availability, and serviceability (RAS) features in the hardware itself, plays such an important role in scalable server deployments. The HA characteristics of the software environment and the RAS features of the hardware platform combine to support what IBM terms "business resilience" on Power Systems — or the ability to continue business processing, even in the event of outages.

PowerHA includes high-availability protection for the IBM AIX, IBM i, and Linux environments on IBM Power Systems. With IBM i 6.1, PowerHA now delivers a hardware disk clustering feature as an alternative to the traditional, software-based, logical replication disaster recovery methods that are widely used on many IBM System i servers. PowerHA on IBM i is similar to IBM's HACMP disk-based clustering product for IBM AIX and Linux. This approach allows restart via a replicated hardware method rather than via logical data mirroring. IDC notes that IBM also offers the IBM iCluster logical replication product, which it gained through its acquisition of the Canadian company DataMirror in 2007. Also, multiple software ISVs, including Vision Solutions (which acquired iTERA and Lakeview Technology in recent years), continue to provide logical replication high-availability software for IBM i workloads (as well as older IBM i5/OS workloads) running on Power Systems.

The IBM i Operating Environment

In April 2008, IBM announced that the IBM i5/OS operating system would be renamed IBM i and that the operating environment would run on the new POWER6 processor–based Power Systems, alongside IBM AIX and Linux. This action ensures that the IBM i operating environment, which has supported business-critical and mission-critical applications for decades, first on AS/400 and recently on System i, is being brought forward to the new generation of POWER6-based servers and blades. The server models in this series — the IBM Power 520 (both 1-socket and 2-socket models), 550, 570, and 595 — and the JS12 and JS22 server blades will all support IBM i applications without change, preserving customers' previous investments in the i5/OS and OS/400 versions of the operating environment.

On a hardware level, IBM will also provide customers with upgrades of existing systems based on POWER5 and POWER5+ processors to POWER6 processors, providing hardware, software, and services to update the systems that run IBM i5/OS to the new Power Systems models running the current releases of IBM i 5.4 or 6.1. IDC notes that the new support for blades requires the latest IBM i 6.1 release.

Importantly, this platformwide action by IBM brings the IBM i into price and functional parity with IBM AIX and Linux as the three primary operating environments hosted in IBM Power Systems. In a highly virtualized IBM Power System, the IBM i application environment will have the ability to grow, as needed, to support ISV packaged applications, as well as custom applications that have been developed expressly for a given IT organization or the industry in which it operates. This means that customers that have relied on i5/OS or OS/400 will be able to bring their applications forward to the new systems, without change. Since IBM i runs on the new POWER6 at the current 5.4 or 6.1 release levels, there is no need to change or recompile the application code as it moves forward from IBM servers based on POWER4 or POWER5 processors.

Business Benefits of This Approach

The IBM i operating environment now has top-tier status on IBM Power Systems. That means that IBM is showing its commitment to the longtime customers of System i and its predecessors and of applications written to run on those servers. Earlier application investments in this platform, which date back to the introduction of the AS/400 servers in 1988, are being brought forward to run on the latest generation of IBM hardware built on POWER6 processors. Because those systems were the first fully 64-bit RISC servers in the IBM product line, the heritage of assistance for 64-bit workloads that leverage that architecture is intact, and strengthened, by moving to POWER6 systems, with PowerVM virtualization and PowerHA availability software.

From a customer perspective, continuity is in place, as well, with the IBM channel partners continuing to promote Power Systems and IBM i — and the full complement of ISV packaged applications that run on IBM i. These include business applications from Oracle (such as J.D. Edwards applications), SAP, Lawson, Infor, and others, as well as IBM WebSphere, IBM Lotus Domino, and the IBM DB2 database. The side-by-side deployments of IBM i, along with IBM AIX and Linux, will aid in datacenter consolidation and workload consolidation efforts that are being undertaken to reduce operational costs. Flexibility in deployment will also provide an opportunity for businesses to make changes in deployment patterns without taking a "rip-and-replace" approach to updating their IT infrastructure. This, too, will help to reduce operational costs and will promote business agility.

New Functionality for IBM i 6.1

The functionality of IBM i release 6.1 shows a continued evolution of its capabilities supporting business-critical applications. Highlights of the latest round of feature announcements include the following:

- **Optimized performance for SAN storage.** With IBM i 6.1 and POWER6-based servers, IBM is changing its storage strategy for large enterprises running IBM i applications to switch from optimizing for "in-the-box" or internally virtualized storage to optimizing for externally virtualized storage centered on a SAN architecture. IBM has changed its approach to fit in with the SAN storage strategy that most large enterprises are implementing today.

  The expected benefits associated with this approach include availability of more storage capacity for virtualized workloads; improved resource utilization on SANs supporting multiple operating systems — including IBM AIX, IBM i, Linux, and Windows — on the same SAN array; and improved support for disaster recovery across multiple systems within the datacenter that can access two or more SANs for faster recovery of workloads in the event of an outage.

- **Same Java Virtual Machine (JVM) for IBM AIX, IBM i, and Linux.** Supporting fully 64-bit Java-based computing environments, the JVM on IBM Power Systems is expected to improve the performance of Java-based workloads running on IBM i, making it comparable to the performance achieved with IBM AIX and Linux. Also, the single JVM across the three operating systems is expected to improve ease of application portability and provide more consistency in Java tools and performance tuning. IDC notes that all three environments also support open-source development for Web-based workloads, including Apache, the open-source MySQL database, and the PHP language (also known as the "AMP" software stack from the widely used "open-source "LAMP" stack).

- **IBM BladeCenter support for IBM i running on POWER6 blades.** IBM is supporting IBM i on blades for the first time with this new release. Two new POWER6 blades are covered: the 2-socket JS12 blade and the 4-socket JS22 blade. Both types of blades can be housed in either an IBM BladeCenter H chassis that holds to 14 blades or an IBM BladeCenter S chassis that supports up to 6 blades with slots for storage built into the chassis. This brings all applications running under the IBM i environment into a bladed server form factor, making the bladed system an integration point for the datacenter and allowing IBM i workloads to run alongside other workloads (e.g., IBM AIX and Linux on POWER processor–based blades and both Linux and Microsoft Windows on x86 blades) for more flexible system management.

- **Support for nondisruptive upgrades from older IBM POWER-based systems to the new line of POWER6-based systems.** IBM provides serial number–protected upgrades from selected POWER5 and POWER5+ models to the new POWER6 processor–based servers. Also, the new servers run either of the two current releases, IBM i 5.4 or 6.1, so no application changes are required to access the new POWER6 servers.

- **Other features** include role-based security that controls access to applications and enablement of micropartitioning on the POWER processors, with 10 micropartitions supported per POWER processor core.

CHALLENGES AND OPPORTUNITIES

Challenges in the Marketplace

The server market continues to be a highly competitive one, with competition from other server vendors in the volume, midrange, and high-end segments of the worldwide server market (including the other top vendors in the worldwide server market — Dell, Fujitsu/Fujitsu Siemens, HP, and Sun Microsystems — listed in alphabetical order). This ongoing competition has the practical effect of reducing system prices, over time, and benefits customers that are evaluating, purchasing, or leasing new technology. On a software level, the use of standards (e.g., Java, Linux, and open source software) has also meant that a wide variety of server products are available to run new workloads in the enterprise, ensuring competition on the basis of price, performance, and technical features.

Meeting the Challenges

By continuing to innovate in its POWER processor architecture, IBM is bringing more powerful solutions to customers, and it is doing so at competitive prices. By supporting multiple computing environments on POWER, IBM is providing both choice and coverage for the range of business applications being used at customer sites today. By supporting Internet protocols and related Web technologies, such as Java, Web interfaces, and the PHP programming language, IBM is hosting a blend of workloads — reflecting the diversity of workloads in today's enterprise (IDC notes that Power Systems do not currently host Windows workloads but that IBM has interconnect features that allow Power Systems to interact with Windows-based servers in the datacenter). By providing advanced functionality in the virtualized environment within the datacenter, IBM i is bringing forward valuable application workloads that can be managed across computing resources on Power Systems, building a more flexible computing environment for customer use.

Opportunities

Business processes, and support for transaction processing applications and LOB applications, have long been the strong suit for IBM System i. This approach is "in sync" with the move toward dynamic IT, which aligns business processes and IT infrastructure more closely, creating an on-demand environment for end users in business units across the enterprise organization.

The flexibility of workload deployment on IBM i platforms, assigning computing tasks to available resources and moving them as needed, speaks to the enterprise needs for business agility, as business requirements change. The IBM i operating environment enables high availability for those business applications, improving uptime of the systems. This impacts operational expenses by improving overall productivity and avoiding IT staff costs related to system maintenance and updates. The IBM i operating environment's strong support for virtualization technologies means that it has the ability to run production-level enterprise applications in a virtualized environment today, without further engineering or application of additional software technologies. Customers considering how to deploy enterprise workloads in a virtualized, production environment have the option of

consolidating workloads onto Power Systems running IBM i as one alternative that would be able to run, in production, with strong support for virtualized I/O and memory bandwidth.

CONCLUSION

Business processes and IT processes are becoming aligned, following years when IT infrastructure was deployed and optimized to meet the needs of separate, and competing, business units across the enterprise. This approach to a more unified virtualization platform is an important step on the road to dynamic IT — mapping business processes to IT infrastructure and enabling the enterprise to support new workloads more seamlessly, as business requirements change, over time.

With Power Systems, and access to local disk or to SANs, workloads running on IBM i have access to the full capacity of scalable servers and networked storage so that rapidly growing workloads can easily gain access to sufficient computing resources, as needed. This addresses the needs of enterprise customers that see demands for applications growing over time and also adds capacity on demand during "crunch times" when peak demands force a decision about "rightsizing" an application on systems that have run out of sufficient "headroom" to add resources quickly and conveniently. Given the changes on both the hardware side and the software side of the platform, IBM can now add more physical server capacity, even as the software components of the solution gain rapid access to more server and storage resources.

These server and operating system offerings are aligned with the major trends in the IT industry: virtualization, consolidation, support for high availability, and support for green IT through efficient use of IT resources. Longtime customers will find their applications and software "stacks" of operating systems and middleware supported on the new Power Systems, and new customers will find that the same systems enable IT initiatives that are making infrastructure more flexible, and more responsive, to ongoing change in the customers' business.